



SUSE Linux Enterprise Server 15 SP6

Deployment Guide

Deployment Guide

SUSE Linux Enterprise Server 15 SP6

This guide details how to install single or multiple systems, and how to exploit the product-inherent capabilities for a deployment infrastructure.

Publication Date: February 27, 2025

<https://documentation.suse.com> 

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

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

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 [xiii](#)

- 1 Available documentation [xiii](#)
- 2 Improving the documentation [xiv](#)
- 3 Documentation conventions [xv](#)
- 4 Support [xvii](#)
Support statement for SUSE Linux Enterprise Server [xvii](#) • Technology
previews [xviii](#)

I [INSTALLATION PREPARATION](#) [1](#)

1 [Planning for SUSE Linux Enterprise Server](#) [2](#)

- 1.1 Considerations for deployment of SUSE Linux Enterprise Server [2](#)
- 1.2 Deployment of SUSE Linux Enterprise Server [3](#)
- 1.3 Running SUSE Linux Enterprise Server [3](#)
- 1.4 Registering SUSE Linux Enterprise Server [4](#)
- 1.5 Changes in installation from SUSE Linux Enterprise Server version
15 [4](#)
Unified Installer for SUSE Linux Enterprise-based products [4](#) • Installing
with Internet access [5](#) • Offline installation [5](#) • Quarterly updated
media [6](#)

2 [Installation on AMD64 and Intel 64](#) [7](#)

- 2.1 Hardware requirements [7](#)
- 2.2 Installation considerations [8](#)
Installation on hardware or virtual machine [9](#) • Installation target [9](#)
- 2.3 Installation methods [9](#)

2.4	Booting the system	10
2.5	Dealing with boot and installation problems	11
	Problems booting	12 • Problems installing 13 • Initiating installation instead of booting 13
3	Installation on AArch64	14
3.1	Hardware requirements	14
3.2	Installation considerations	15
	Installation on hardware or virtual machine	16 • Installation target 16
3.3	Controlling the installation process	16
3.4	Booting the system	17
3.5	Dealing with boot and installation problems	18
	Boot problems	19 • Problems installing 19 • Initiating installation instead of booting 20
3.6	Raspberry Pi	20
	Boot process	20 • Lack of a real-time clock 23 • Deploying an appliance image 23 • Installation from USB media 24 • Installation from network 25 • More information 25
4	Installation on IBM POWER	26
4.1	Hardware requirements	26
4.2	Installing SUSE Linux Enterprise Server for POWER	27
4.3	Installing SUSE Linux Enterprise Server	30
4.4	More information	34
5	Installation on IBM Z and LinuxONE	35
5.1	System requirements	35
	Hardware	35 • MicroCode Level, APARs, and fixes 37 • Software 38
5.2	General information	39
	Installation types	39 • IPL options 40

- 5.3 Preparing for installation 41
 - Making the installation data available 41 • Installation types 49 • Preparing the IPL of the SUSE Linux Enterprise Server installation system 50 • IPLing the SUSE Linux Enterprise Server installation system 53 • Network configuration 58 • Connecting to the SUSE Linux Enterprise Server installation system 62 • The SUSE Linux Enterprise Server boot procedure on IBM Z 64
- 5.4 Secure boot 65
- 5.5 The parmfile—automating the system configuration 66
 - General parameters 67 • Configuring the network interface 68 • Specifying the installation source and YaST interface 70 • I/O device auto-configuration on IBM Z systems 71 • Example parmfiles 72
- 5.6 Using the vt220 terminal emulator 73
- 5.7 More information 74
 - General documents about Linux on IBM Z 74 • Technical issues of Linux on IBM Z 74 • Advanced configurations for Linux on IBM Z 75

6 Installation on virtualization hosts 76

7 Installation on hardware not supported at release 77

- 7.1 Download kernel update 77
- 7.2 Boot kernel update 77

II INSTALLATION PROCEDURE 78

8 Boot parameters 79

- 8.1 Using the default boot parameters 79
- 8.2 PC (AMD64/Intel 64/AArch64) 80
 - The boot screen on machines with traditional BIOS 80 • The boot screen on machines equipped with UEFI 82

8.3	List of important boot parameters	85
	General boot parameters	85 • Configuring the network interface 86 • Specifying the installation source 88 • Specifying remote access 89
8.4	Advanced setups	89
	Providing data to access a Repository Mirroring Tool server	90 • Configuring an alternative data server for supportconfig 91 • Using IPv6 for the installation 91 • Using a proxy for the installation 91 • Enabling SELinux support 92 • Enabling the installer self-update 93 • Reusing LVM 93 • Scale user interface for high DPI 93 • Using CPU mitigations 93 • LUKS 2 Support 94
8.5	IBM Z	94
8.6	More information	96
9	Installation steps	97
9.1	Overview	97
9.2	Installer self-update	98
	Self-update process	99 • Custom self-update repositories 101
9.3	Language, keyboard and product selection	103
9.4	License agreement	105
9.5	IBM Z: disk activation	105
	Configuring DASD disks	106 • Configuring zFCP disks 107
9.6	Network settings	109
9.7	Registration	110
	Manual registration	111 • Loading registration codes from USB storage 112 • Installing without registration 113
9.8	Extension and module selection	115
9.9	Add-on product	120
9.10	System roles	121

- 9.11 Partitioning 123
 - Important information 123 • Suggested partitioning 126
- 9.12 Clock and time zone 128
- 9.13 Create new user 129
- 9.14 Authentication for the system administrator root 132
- 9.15 Installation settings 134
 - Software 135 • Booting 137 • Security 137 • Security Profiles 139 • Network configuration 139 • Kdump 140 • IBM Z: blacklist devices 140 • Default systemd target 140 • Import SSH host keys and configuration 141 • System 141*
- 9.16 Performing the installation 142
 - IBM Z: IPLing the installed system 142 • IBM Z: Connecting to the installed system 144

10 Registering SUSE Linux Enterprise and managing modules/extensions 146

- 10.1 Registering during the installation 147
- 10.2 Registering during automated deployment 147
- 10.3 Registering from the installed system 147
 - Registering with YaST 147 • Registering with SUSEConnect 150
- 10.4 Managing modules and extensions in a running system 151
 - Adding modules and extensions with YaST 151 • Deleting modules and extensions with YaST 152 • Adding or deleting modules and extensions with SUSEConnect 153
- 10.5 SUSEConnect keep-alive timer 155

11 *Expert Partitioner* 157

- 11.1 Using the *Expert Partitioner* 157
 - Partition tables 159 • Partitions 160 • Editing a partition 164 • Expert options 167 • Advanced options 167 • More partitioning tips 167 • Partitioning and LVM 170

- 11.2 Device encryption **170**
 - Encryption methods **171** • Password-based key derivation functions **173**
- 11.3 LVM configuration **174**
 - Create physical volume **174** • Creating volume groups **175** • Configuring logical volumes **176**
- 11.4 Soft RAID **178**
 - Soft RAID configuration **178** • Troubleshooting **179** • More information **180**

12 Remote installation **181**

- 12.1 Overview **181**
- 12.2 Scenarios for remote installation **182**
 - Installation from source media via VNC **182** • Network installation using VNC **183** • Installation from source media via SSH **184** • Installation from network via SSH **185**
- 12.3 Monitoring installation via VNC **185**
 - Preparing for VNC installation **186** • Connecting to the installation program **186**
- 12.4 Monitoring installation via SSH **187**
 - Preparing for SSH installation **187** • Connecting to the installation program **187**
- 12.5 Installation via serial console **188**

13 Troubleshooting **189**

- 13.1 Checking media **189**
- 13.2 No bootable drive available **189**
- 13.3 Booting from installation media fails **190**
- 13.4 Boot failure **191**
- 13.5 Graphical installer fails to start **193**
- 13.6 Only minimal boot screen is displayed **194**

13.7 Log files 195

III CUSTOMIZING INSTALLATION IMAGES 196

14 Prepare a disk for cloning with the system cleanup tool 197

14.1 Cleaning up unique system identifiers 197

15 Customizing installation images with mksusecd 198

15.1 Installing mksusecd 198

15.2 Creating a minimal boot image 199

15.3 Setting default kernel boot parameters 199

15.4 Customizing modules, extensions, and repositories 199

15.5 Creating a minimal netinstall ISO 201

15.6 Change default repository 201

16 Customizing installation images manually 202

IV SETTING UP AN INSTALLATION SERVER 203

17 Setting up a network installation source 204

17.1 Setting up an installation server using YaST 204

17.2 Setting up an NFS repository manually 206

17.3 Setting up an FTP repository manually 209

17.4 Setting up an HTTP repository manually 210

17.5 Managing an SMB repository 211

17.6 Using ISO images of the installation media on the server 212

18 Preparing network boot environment 214

- 18.1 Setting up a DHCP server 214
 - Dynamic address assignment 215 • Assigning static IP addresses 216 • PXE and AutoYaST installation failures 216
- 18.2 Setting up a TFTP server 217
 - Installing a TFTP server 217 • Installing files for boot 218 • Configuring PXELINUX 219 • Preparing PXE boot for EFI with GRUB2 220
- 18.3 PXELINUX configuration options 220
- 18.4 Preparing the target system for PXE boot 223
- 18.5 Using wake-on-LAN for remote wakeups 223
 - Prerequisites 223 • Verifying wired Ethernet support 223 • Verifying wireless interface support 224 • Installing and testing WOL 225

19 Setting up a UEFI HTTP Boot server 226

- 19.1 Introduction 226
 - Configuring the client machine 226 • Preparation 226
- 19.2 Configuring the server 226
 - DNS server 226
- 19.3 Booting the client via HTTP boot 234

20 Deploying customized preinstallations 235

- 20.1 Preparing the master machine 235
- 20.2 Customizing the firstboot installation 236
 - Customizing YaST messages 237 • Customizing the license action 238 • Customizing the release notes 238 • Customizing the workflow 239 • Configuring additional scripts 243 • Providing translations of the installation workflow 243
- 20.3 Cloning the master installation 244
- 20.4 Personalizing the installation 244

- A Imaging and creating products 245
- B GNU licenses 246

Preface

1 Available documentation

Online documentation

Our documentation is available online at <https://documentation.suse.com>. Browse or download the documentation in various formats.



Note: Latest updates

The latest updates are usually available in the English-language version of this documentation.

SUSE Knowledgebase

If you run into an issue, 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/releasesnotes/>.

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 end of the line:

```
> 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 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

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 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/releasesnotes> .

I Installation preparation



- 1 Planning for SUSE Linux Enterprise Server 2
- 2 Installation on AMD64 and Intel 64 7
- 3 Installation on AArch64 14
- 4 Installation on IBM POWER 26
- 5 Installation on IBM Z and LinuxONE 35
- 6 Installation on virtualization hosts 76
- 7 Installation on hardware not supported at release 77

1 Planning for SUSE Linux Enterprise Server

This chapter describes some basic considerations before installing SUSE Linux Enterprise Server.

1.1 Considerations for deployment of SUSE Linux Enterprise Server

The implementation of an operating system either in an existing IT environment or as a completely new rollout must be carefully prepared. At the beginning of the planning process, you should try to define the project goals and necessary features. This must always be done individually for each project, but the questions to answer should include the following:

- How many installations should be done? Depending on this, the best deployment methods differ.
- Will the system run as physical host or as a virtual machine?
- Will the system be exposed to external threats like hacker attacks? Have a look at *Book "Security and Hardening Guide", Chapter 1 "Security and confidentiality"* to get an overview of consequences.
- How will you get regular updates? All patches are provided online for registered users in the [SUSE Customer Center \(https://scc.suse.com/\)](https://scc.suse.com/) .
- Do you need help for your local installation? SUSE provides training, support, and consulting for all topics pertaining to SUSE Linux Enterprise Server. Find more information about this at <https://www.suse.com/products/server/> .
- Do you need third-party products? Make sure that the required product is also supported on the desired platform. SUSE can provide help to support software on different platforms when needed.

1.2 Deployment of SUSE Linux Enterprise Server

To make sure that your system will run flawlessly, always try to use certified hardware. The hardware certification process is an ongoing process and the database of certified hardware is updated regularly. Find the search form for certified hardware at <https://www.suse.com/yessearch/Search.jsp>.

Depending on the number of desired installations, it is beneficial to use installation servers or even completely automatic installations. When using Xen or KVM virtualization technologies, network root file systems or network storage solutions like iSCSI should be considered.

SUSE Linux Enterprise Server provides you with a broad variety of services. Most of the needed configurations can be made with YaST, the SUSE configuration utility. In addition, many manual configurations are described in the corresponding chapters.

In addition to the plain software installation, you should consider training the end users of the systems and help desk staff.



Note: Terminology

In the following sections, the system to hold your new SUSE Linux Enterprise Server installation is called *target system* or *installation target*. The term *repository* (previously called “installation source”) is used for all sources of installation data. This includes physical media, such as CD, DVD, or USB flash drive, and network servers distributing the installation data in your network.


1.3 Running SUSE Linux Enterprise Server

The SUSE Linux Enterprise Server operating system is a well-tested and stable system. Unfortunately, this does not prevent hardware failures or other causes for downtime or data loss. Make sure that you have a backup solution in place for mission-critical tasks.

For optimal security and data safety, you should make regular updates of all the operated machines. If you have a mission critical server, you should run a second identical (pre-production) machine that you can use to test all changes. This also gives you the possibility of switching machines in case of hardware failure.

1.4 Registering SUSE Linux Enterprise Server

To get technical support and product updates, you need to register and activate your SUSE product with the SUSE Customer Center. We recommend to register during the installation, since this will enable you to install the system with the latest updates and patches available. However, if you are offline or want to skip the registration step, you can complete registration from the running system.

In case your organization does not provide a local registration server, registering SUSE Linux Enterprise requires a SUSE Customer Center account. In case you do not have one yet, go to the SUSE Customer Center home page (<https://scc.suse.com/> ) to create one.

During the installation you will be asked to enter your registration code. For details, see [Section 9.7, “Registration”](#).

If you deploy your instances automatically using AutoYaST, you can register the system during the installation by providing the respective information in the AutoYaST control file. For details, see *Book “AutoYaST Guide”, Chapter 4 “Configuration and installation options”, Section 4.3 “System registration and extension selection”*.

For registering an already installed system, see *Book “Administration Guide”, Chapter 8 “Installing or removing software”, Section 8.2 “Registering an installed system”*.

1.5 Changes in installation from SUSE Linux Enterprise Server version 15

Starting with SUSE Linux Enterprise Server 15, all SUSE Linux Enterprise-based products are installed using a Unified Installer from a single set of installation media for each supported architecture.

1.5.1 Unified Installer for SUSE Linux Enterprise-based products

As of SUSE Linux Enterprise Server 15 SP6, this includes the following base products:

Product Name	Supported Platforms
SUSE Linux Enterprise Server	AMD64/Intel 64; AArch64; POWER; IBM Z

Product Name	Supported Platforms
SUSE Linux Enterprise High Performance Computing	AMD64/Intel 64; AArch64
SUSE Linux Enterprise Real Time	AMD64/Intel 64
SUSE Linux Enterprise Server for SAP Applications	AMD64/Intel 64; POWER
SUSE Linux Enterprise Desktop	AMD64/Intel 64
SUSE Manager Server	AMD64/Intel 64; POWER; IBM Z
SUSE Manager Proxy	AMD64/Intel 64
SUSE Manager for Retail Branch Server	AMD64/Intel 64
SUSE Enterprise Storage	AMD64/Intel 64; Arm; Intel 64

1.5.2 Installing with Internet access

If you are installing onto a computer or VM that has access to the Internet, then to install any of the products listed above, it is only necessary to download the [SLE-15-SP6-Online-ARCH-GM-media1.iso](#) image for the desired architecture.



Note: Installing SUSE Manager

To install any SUSE Manager products, the target machine must have direct access to the SUSE Customer Center or to an RMT server.

1.5.3 Offline installation

Except for SUSE Manager, you do not require access to the Internet, or to the SUSE Customer Center or to an Repository Mirroring Tool server, to install the other listed products.

For offline installation, additionally download the [SLE-15-SP6-Full-ARCH-GM-media1.iso](#) image for the desired architecture.

There is an additional, second Packages medium, but this contains only source code and is not required for installation.



Tip: Full media size

The size of the full installation media SLE-15-SP6-Online-ARCH-GM-media1.iso exceeds the capacity of a dual layer DVD. Therefore you can only boot it from a USB flash drive.

1.5.4 Quarterly updated media

For the installation media and the VM Guest images, SUSE offers two variants:

- The first, containing GM in the file name, consists of the package set as shipped on the first customer shipment date.
- The second, identified by a QU followed by a number in the file name, contains the same package set but includes all maintenance updates of the packages that have been released in the meantime. Quarterly updated media are refreshed every three months, with the first coming three months after the GM release.

You only need either the GM or the QU media, not both. Which version to select depends on your needs and preferences. If you have newer hardware, the QU version might be the better choice. The installation procedure is identical for both variants.

For both variants it is recommended to install the latest updates released after creation of the images during or immediately after installation.

2 Installation on AMD64 and Intel 64

This chapter describes the steps necessary to prepare for the installation of SUSE Linux Enterprise Server on AMD64 and Intel 64 computers. It introduces the steps required to prepare for various installation methods. The list of hardware requirements provides an overview of systems supported by SUSE Linux Enterprise Server. Find information about available installation methods and several commonly known problems. Also learn how to control the installation, provide installation media, and boot with regular methods.

2.1 Hardware requirements

The SUSE® Linux Enterprise Server operating system can be deployed on a wide range of hardware. It is impossible to list all the different combinations of hardware SUSE Linux Enterprise Server supports. However, to provide you with a guide to help you during the planning phase, the minimum requirements are presented here.

If you want to be sure that a given computer configuration will work, find out which platforms have been certified by SUSE. Find a list at <https://www.suse.com/yessearch/>.

CPU

Most CPUs available at the time of release are supported.

Maximum number of CPUs

The maximum number of CPUs supported by software design is 8192 for Intel 64 and AMD64. If you plan to use such a large system, verify with our hardware system certification Web page for supported devices, see <https://www.suse.com/yessearch/>.

Memory requirements

A minimum of 1024 MB of memory is required for a minimal installation. On machines with more than two processors, add 512 MB per CPU. For remote installations via HTTP or FTP, add another 150 MB. Note that these values are only valid for the installation of the operating system—the actual memory requirement in production depends on the system's workload. For systems running the GNOME desktop environment, a minimum of 2048 MB of memory is required and 4096 MB is recommended.

Hard disk requirements

The disk requirements depend largely on the installation selected and how you use your machine. Commonly, you need more space than the installation software itself needs to have a system that works properly. Minimum requirements for different selections are:

Installation Scope	Minimum Hard Disk Requirements
Text Mode	1.5 GB
Minimal System	2.5 GB
GNOME Desktop	3 GB
All patterns	4 GB
Recommended Minimum (no Btrfs snapshots): 10 GB	
Required Minimum (with Btrfs snapshots): 16 GB	
Recommended Minimum (with Btrfs snapshots): 32 GB	

If your root partition is smaller than 10 GB, the installer will not make an automated partitioning proposal and you need to manually create partitions. Therefore the recommended minimum size for the root partition is 10 GB. If you want to enable Btrfs snapshots on the root volume to enable system rollbacks (see *Book “Administration Guide”, Chapter 10 “System recovery and snapshot management with Snapper”*) the minimum size for the root partition is 16 GB.

Boot methods

The computer can be booted from a CD or a network. A special boot server is required to boot over the network. This can be set up with SUSE Linux Enterprise Server.

2.2 Installation considerations

This section encompasses many factors that need to be considered before installing SUSE Linux Enterprise Server on AMD64 and Intel 64 hardware.

2.2.1 Installation on hardware or virtual machine

SUSE Linux Enterprise Server is normally installed as an independent operating system. With virtualization it is also possible to run multiple instances of SUSE Linux Enterprise Server on the same hardware. However, the installation of the VM Host Server is performed like a typical installation with some additional packages. The installation of virtual guests is described in *Book "Virtualization Guide", Chapter 10 "Guest installation"*.

2.2.2 Installation target

Most installations are to a local hard disk. Therefore, it is necessary for the hard disk controllers to be available to the installation system. If a special controller (like a RAID controller) needs an extra kernel module, provide a kernel module update disk to the installation system.

Other installation targets may be various types of block devices that provide sufficient disk space and speed to run an operating system. This includes network block devices like iSCSI or SAN. It is also possible to install on network file systems that offer the standard Unix permissions. However, it may be problematic to boot these, because they must be supported by the initramfs before the actual system can start. Such installations can be useful when you need to start the same system in different locations or you plan to use virtualization features like domain migration.

2.3 Installation methods

You can choose the desired installation method by booting the setup with one of the options listed in [Section 2.4, "Booting the system"](#). To enable the additional installation methods, refer to [Section 8.3.4, "Specifying remote access"](#). For information about how to use remote installation methods, refer to [Chapter 12, Remote installation](#).

A brief overview of the different methods:

Local with monitor and keyboard

This is the method most frequently used to install SUSE Linux Enterprise Server. This also requires very little preparation but needs a lot of direct interaction.

Remote via SSH

You can perform installation via SSH either in text mode or use X-forwarding for a graphical installation. For details, refer to [Section 12.4, "Monitoring installation via SSH"](#).

Remote via serial console

For this installation method, you need a second computer connected via a *null modem* cable to the target computer. The installation is done in text mode. For details, refer to [Section 12.5, “Installation via serial console”](#).

Remote via VNC

Use this method to perform the installation using a graphical interface without direct access to the target machine. For details, refer to [Section 12.3, “Monitoring installation via VNC”](#).

Automatic via AutoYaST

To install SUSE Linux Enterprise Server on several computers with similar hardware, it is recommended you perform the installation using AutoYaST. In this case, start by installing one SUSE Linux Enterprise Server and use it to create the necessary AutoYaST configuration files. For details, refer to *Book “AutoYaST Guide”*.

2.4 Booting the system

This section gives an overview of the steps required for the complete installation of SUSE® Linux Enterprise Server.

Unlike previous SLE products, the entire SLE 15 SP6 product line can be installed using the Unified Installer. For details about the changes since SUSE Linux Enterprise 15 and which media to download for installation, see [Section 1.5, “Changes in installation from SUSE Linux Enterprise Server version 15”](#).

For a full description of how to install and configure the system with YaST, refer to [Part II, “Installation procedure”](#).



Important: Hardware support updates

When using very recent hardware, it can be necessary to boot the installation with a newer kernel from a [Kernel Update ISO](#) image. For details, refer to [Chapter 7, Installation on hardware not supported at release](#).

1. Prepare the installation media.

USB Flash Drive

This is the simplest way to start the installation. To create a bootable flash disk, you need to copy a DVD image to the device using the `dd` command. The flash disk must not be mounted, and all data on the device will be erased.

```
# dd if=PATH_TO_ISO_IMAGE of=USB_STORAGE_DEVICE bs=4M
```

Network booting

If the target computer's firmware supports it, you can boot the computer from the network and install from a server. This booting method requires a boot server that provides the needed boot images over the network. The exact protocol depends on your hardware. Commonly you need several services, such as TFTP and DHCP or PXE boot. For details, read [Chapter 18, Preparing network boot environment](#).

It is possible to install from many common network protocols, such as NFS, HTTP, FTP, or SMB. For more information on how to perform such an installation, refer to [Chapter 12, Remote installation](#).

2. Configure the target system firmware to boot the medium you chose. Refer to the documentation of your hardware vendor about how to configure the correct boot order.
3. Set the boot parameters required for your installation control method. An overview of the different methods is provided in [Section 2.3, "Installation methods"](#). A list of boot parameters is available in [Chapter 8, Boot parameters](#).
4. Perform the installation as described in [Chapter 9, Installation steps](#). The system needs to restart after the installation is finished.
5. Optional: Change the boot order of the system to directly boot from the medium to which SUSE Linux Enterprise Server has been installed. If the system boots from the installation medium, the first boot parameter will be to boot the installed system.

2.5 Dealing with boot and installation problems

Prior to delivery, SUSE® Linux Enterprise Server is subjected to an extensive test program. Despite this, problems occasionally occur during boot or installation.

2.5.1 Problems booting

Boot problems may prevent the YaST installer from starting on your system. Another symptom is when your system does not boot after the installation has been completed.

System does not boot from installation media

Change your computer's firmware or BIOS so that the boot sequence is correct. To do this, consult the manual for your hardware.

The computer hangs

Change the console on your computer so that the kernel outputs are visible. Be sure to check the last outputs. This is normally done by pressing **Ctrl – Alt – F10** . If you cannot resolve the problem, consult the SUSE Linux Enterprise Server support staff. To log all system messages at boot time, use a serial connection as described in [Section 2.3, “Installation methods”](#).

Boot disk

The boot disk is a useful interim solution if you have difficulties setting the other configurations or if you want to postpone the decision regarding the final boot mechanism. For more details on creating boot disks, see *Book “Administration Guide”, Chapter 18 “The boot loader GRUB 2”* `grub2-mkrescue`.

Virus warning after installation

There are BIOS variants that check the structure of the boot sector (MBR) and erroneously display a virus warning after the installation of GRUB 2. Solve this problem by entering the BIOS and looking for corresponding adjustable settings. For example, switch off *virus protection*. You can switch this option back on again later. It is unnecessary, however, if Linux is the only operating system you use.

2.5.2 Problems installing

If an unexpected problem occurs during installation, information is needed to determine the cause of the problem. Use the following directions to help with troubleshooting:

- Check the outputs on the various consoles. You can switch consoles with the key combination `Ctrl – Alt – Fn`. For example, obtain a shell in which to execute various commands by pressing `Ctrl – Alt – F2`.
- Try launching the installation with “Safe Settings” (press `F5` on the installation screen and choose *Safe Settings*). If the installation works without problems in this case, there is an incompatibility that causes either ACPI or APIC to fail. In some cases, a BIOS or firmware update fixes this problem.
- Check the system messages on a console in the installation system by entering the command `dmesg -T`.

2.5.3 Initiating installation instead of booting

The default option in the boot menu of the installation source for SUSE Linux Enterprise Server boots the machine into the already installed system. To avoid this and to initiate the installation process instead, choose one of the available installation options in the boot menu.

3 Installation on AArch64

This chapter describes the steps necessary to prepare for the installation of SUSE Linux Enterprise Server on AArch64 computers. It introduces the steps required to prepare for various installation methods. The list of hardware requirements provides an overview of systems supported by SUSE Linux Enterprise Server. Find information about available installation methods and several common known problems. Also learn how to control the installation, provide installation media, and boot with regular methods.

3.1 Hardware requirements

The SUSE® Linux Enterprise Server operating system can be deployed on a wide range of hardware. It is impossible to list all the different combinations of hardware SUSE Linux Enterprise Server supports. However, to provide you with a guide to help you during the planning phase, the minimum requirements are presented here.

If you want to be sure that a given computer configuration will work, find out which platforms have been certified by SUSE. Find a list at <https://www.suse.com/yessearch/>.

CPU

The minimum requirement is a CPU that supports the Armv8-A instruction set architecture (ISA), for example, Arm Cortex-A53 or Cortex-A57. Refer to <https://www.arm.com/products/processors/cortex-a/> for a list of available Armv8-A processors.

CPUs with the Armv8-R (realtime) and Armv8-M (microcontroller) ISA are currently not supported.

Maximum number of CPUs

The maximum number of supported CPUs is 256. If you plan to use such a large system, check our hardware system certification Web page for supported devices, see <https://www.suse.com/yessearch/>.

Memory requirements

A minimum of 1024 MB of memory is required for a minimal installation. On machines with more than two processors, add 512 MB per CPU. For remote installations via HTTP or FTP, add another 150 MB. Note that these values are only valid for the installation of the

operating system—the actual memory requirement in production depends on the system's workload. For systems running the GNOME desktop environment, a minimum of 2048 MB of memory is required and 4096 MB is recommended.

Hard disk requirements

The disk requirements depend largely on the installation selected and how you use your machine. Commonly, you need more space than the installation software itself needs to have a system that works properly. Minimum requirements for different selections are:

Installation Scope	Minimum Hard Disk Requirements
Text Mode	1.5 GB
Minimal System	2.5 GB
GNOME Desktop	3 GB
All patterns	4 GB
Recommended Minimum (no Btrfs snapshots): 10 GB	
Required Minimum (with Btrfs snapshots): 16 GB	
Recommended Minimum (with Btrfs snapshots): 32 GB	

If your root partition is smaller than 10 GB, the installer will not make an automated partitioning proposal and you need to manually create partitions. Therefore the recommended minimum size for the root partition is 10 GB. If you want to enable Btrfs snapshots on the root volume to enable system rollbacks (see *Book “Administration Guide”, Chapter 10 “System recovery and snapshot management with Snapper”*) the minimum size for the root partition is 16 GB.

Boot methods

The computer can be booted from a USB disk or a network. A special boot server is required to boot over the network. This can be set up with SUSE Linux Enterprise Server.

3.2 Installation considerations

This section encompasses many factors that need to be considered before installing SUSE Linux Enterprise Server on AArch64 hardware.

3.2.1 Installation on hardware or virtual machine

SUSE Linux Enterprise Server is normally installed as an independent operating system. With virtualization it is also possible to run multiple instances of SUSE Linux Enterprise Server on the same hardware. The installation of the VM Host Server is performed like a typical installation with some additional packages. The installation of virtual guests is described in *Book "Virtualization Guide", Chapter 10 "Guest installation"*.

3.2.2 Installation target

Most installations are to a local hard disk. Therefore, it is necessary for the hard disk controllers to be available to the installation system. If a special controller (like a RAID controller) needs an extra kernel module, provide a kernel module update disk to the installation system.

Other installation targets may be various types of block devices that provide sufficient disk space and speed to run an operating system. This includes network block devices like iSCSI or SAN. It is also possible to install on network file systems that offer the standard Unix permissions. However, it may be problematic to boot these, because they must be supported by the initramfs before the actual system can start. Such installations can be useful when you need to start the same system in different locations or you plan to use virtualization features like domain migration.

3.3 Controlling the installation process

You can choose the desired installation method by booting the setup with one of the options listed in [Section 2.4, "Booting the system"](#). To enable the additional installation methods, refer to [Section 8.3.4, "Specifying remote access"](#). For information about how to use remote installation methods, refer to [Chapter 12, Remote installation](#).

A brief overview of the different methods:

Local with monitor and keyboard

This is the method most frequently used to install SUSE Linux Enterprise Server. This also requires little preparation but needs a lot of direct interaction.

Remote via SSH

You can perform installation via SSH either in text mode or use X-forwarding for a graphical installation. For details, refer to [Section 12.4, "Monitoring installation via SSH"](#).

Remote via serial console

For this installation method, you need a second computer connected via a *null modem* cable to the target computer. The installation is done in text mode. For details, refer to [Section 12.5, “Installation via serial console”](#).

Remote via VNC

Use this method to perform the installation using a graphical interface without direct access to the target machine. For details, refer to [Section 12.3, “Monitoring installation via VNC”](#).

Automatic via AutoYaST

To install SUSE Linux Enterprise Server on several computers with similar hardware, it is recommended you perform the installation using AutoYaST. In this case, start by installing one SUSE Linux Enterprise Server and use it to create the necessary AutoYaST configuration files. For details, refer to *Book “AutoYaST Guide”*.

3.4 Booting the system

This section gives an overview of the steps required for the complete installation of SUSE® Linux Enterprise Server.

Unlike previous SLE products, the entire SLE 15 SP6 product line can be installed using the Unified Installer. For details about the changes since SUSE Linux Enterprise 15 and which media to download for installation, see [Section 1.5, “Changes in installation from SUSE Linux Enterprise Server version 15”](#).

For a full description of how to install and configure the system with YaST, refer to [Part II, “Installation procedure”](#).



Important: Hardware support updates

When using recent hardware, it can be necessary to boot the system with a newer kernel from a [Kernel Update ISO](#) image. For details, refer to [Chapter 7, Installation on hardware not supported at release](#).

1. Prepare the installation media.

USB Flash Drive

This is the simplest way to start the installation. To create a bootable flash disk, you need to copy a DVD image to the device using the `dd` command. The flash disk must not be mounted, and all data on the device will be erased.

```
# dd if=PATH_TO_ISO_IMAGE of=USB_STORAGE_DEVICE bs=4M
```

Network booting

If the target computer's firmware supports it, you can boot the computer from the network and install from a server. This booting method requires a boot server that provides the needed boot images over the network. The exact protocol depends on your hardware. Commonly you need several services, such as TFTP and DHCP or PXE boot. For details, read [Chapter 18, Preparing network boot environment](#).

It is possible to install from many common network protocols, such as NFS, HTTP, FTP, or SMB. For more information on how to perform such an installation, refer to [Chapter 12, Remote installation](#).

2. Configure the target system firmware to boot the medium you chose. Refer to the documentation of your hardware vendor about how to configure the correct boot order.
3. Set the boot parameters required for your installation control method. An overview of the different methods is provided in [Section 3.3, "Controlling the installation process"](#). A list of boot parameters is available in [Chapter 8, Boot parameters](#).
4. Perform the installation as described in [Chapter 9, Installation steps](#). The system needs to restart after the installation is finished.
5. Optional: Change the boot order of the system to directly boot from the medium to which SUSE Linux Enterprise Server has been installed. If the system boots from the installation medium, the first boot parameter will be to boot the installed system.

3.5 Dealing with boot and installation problems

Although SUSE® Linux Enterprise Server undergoes an extensive test program, problems may occasionally occur during boot or installation.

3.5.1 Boot problems

Boot problems may prevent the YaST installer from starting on your system. Another symptom is failure to boot after the installation has been completed.

Machine boots the installed system instead of the installation medium

Change the boot sequence in your machine's BIOS. Refer to the documentation supplied with your hardware for further information.

The system hangs

Change the console on your system so that the kernel outputs are visible. Be sure to check the last few lines of output. This is normally done by pressing `Ctrl – Alt – F10`. If you cannot resolve the problem, consult the SUSE Linux Enterprise Server support staff. To log all system messages at boot time, use a serial connection as described in [Section 2.3, “Installation methods”](#).

Boot disk

The boot disk is a useful interim solution for boot issues. If you have difficulties setting the other configurations, or if you want to postpone the decision regarding the final boot mechanism, use a boot disk. For more details on creating boot disks, see *Book “Administration Guide”, Chapter 18 “The boot loader GRUB 2” grub2-mkrescue*.

3.5.2 Problems installing

If an unexpected problem occurs during installation, information is needed to determine the cause of the problem. Use the following directions to help with troubleshooting:

- Check the outputs on the various consoles. You can switch consoles with the key combination `Ctrl – Alt – Fn`. For example, obtain a shell in which to execute various commands by pressing `Ctrl – Alt – F2`.
- Try launching the installation with “Safe Settings” (press `F5` on the installation screen and choose *Safe Settings*). If the installation works without problems in this case, there is an incompatibility that causes either ACPI or APIC to fail. In some cases, a firmware update fixes this problem.
- Check the system messages on a console in the installation system by entering the command `dmesg -T`.

3.5.3 Initiating installation instead of booting

The default option in the boot menu of the installation medium for SUSE Linux Enterprise Server boots the machine into the already installed system. To initiate the installation process instead, choose one of the available installation options in the boot menu.

3.6 Raspberry Pi

SUSE® Linux Enterprise Server is the first enterprise Linux distribution to support the inexpensive Raspberry Pi* single-board computer. SUSE Linux Enterprise Server 15 SP6 supports the following models:

- Raspberry Pi 3 Model A +
- Raspberry Pi 3 Model B
- Raspberry Pi 3 Model B +
- Raspberry Pi 4 Model B
- Raspberry Pi Compute Module 3
- Raspberry Pi Compute Module 3 +

The Raspberry Pi differs from more conventional server machines in several ways. First and foremost, it does not come with a boot loader capable of loading operating systems. SUSE Linux Enterprise Server therefore ships additional boot loader software to fill that gap.

3.6.1 Boot process

The primary processor on the Raspberry Pi's System-on-Chip (SoC) is the Broadcom VideoCore Graphics Processing Unit (GPU), not the Arm Central Processing Unit (CPU). It is the GPU which starts initializing the hardware from a first-stage boot loader in the on-chip Boot Read-Only Memory (Boot ROM). Only a few configuration options can influence the Boot ROM; see [Section 3.6.1.2, "OTP memory"](#).

The Raspberry Pi 3 hardware does not have any built-in firmware. Instead, its second-stage boot loader firmware `bootcode.bin` is loaded from the boot medium every time the machine is powered on. It in turn loads the third-stage boot loader `start.elf`.

The Raspberry Pi 4 hardware has a small Electrically Erasable Programmable Read-Only Memory (EEPROM) for the second-stage boot loader. Apart from that, its boot sequence is similar to that of the Raspberry Pi 3, loading the third-stage boot loader start4.elf from the boot medium.



Warning: EEPROM update on Raspberry Pi 4

An update of the second-stage boot loader can be performed by booting from a specially prepared microSD card.

Only insert boot media that you trust, and verify that no file called recovery.bin is unintentionally present.

If an armstub8.bin file is present, it will be loaded as a fourth-stage boot loader at AArch64 Exception Level 3 (EL3). Otherwise, a minimal integrated stub will be used.



Note: EL3 security considerations

Code loaded for EL3 (often called BL31) will reside in memory, and Linux may attempt hypercalls into EL3 throughout its runtime.

Verify that your boot media have no armstub8.bin file unintentionally present. SUSE Linux Enterprise Server 15 SP6 does not include it.

Beware that the Raspberry Pi's SoC does not provide TrustZone secure memory. Both the OS on the CPU and any software on the GPU may access its RAM. It is therefore unsuited for cryptographic EL0-s applications. SUSE Linux Enterprise Server does not provide an EL1-s Trusted Execution Environment (TEE) for that reason.

SUSE Linux Enterprise Server for the Raspberry Pi is configured to load a fifth-stage boot loader called Das U-Boot.

3.6.1.1 Config.txt

There is no non-volatile memory to hold configuration information. This means there are no conventional settings to adjust for boot device order, time and date, and so on.

Instead, the boot loader reads a configuration file `config.txt` from the boot medium. The `config.txt` provided by SUSE should not be modified. It allows the user to optionally provide an `extraconfig.txt` file, which can override any setting from `config.txt` if needed. This permits SUSE Linux Enterprise Server to update the `config.txt` file when needed, without overwriting any user settings.

3.6.1.2 OTP memory

The SoC also has a very small amount of One-Time Programmable Memory (OTP memory). This can be used to configure some settings, such as whether the Boot ROM should attempt to boot from USB devices or over Ethernet.

This OTP memory is described on the Raspberry Pi Foundation Web site: <https://www.raspberrypi.org/documentation/hardware/raspberrypi/otpbits.md> ↗



Warning: One-time programmable only

Configuration settings written into OTP memory cannot be reversed.

The most common use case for OTP memory will be enabling USB boot on Raspberry Pi 3 Model B or Compute Module 3.

3.6.1.3 Enabling USB boot mode for Raspberry Pi 3 Model B

To permanently allow booting from connected USB mass storage devices on Raspberry Pi 3 Model B, and from its on-board USB Ethernet, prepare a microSD card as described in [Section 3.6.3, “Deploying an appliance image”](#). Before unmounting or ejecting the card and booting from it, add to its FAT partition a text file `extraconfig.txt` ([Section 3.6.1.1, “Config.txt”](#)) with the following setting:

```
program_usb_boot_mode=1
```

Then continue to boot from the modified microSD card as usual. Once you see output from the U-Boot or GRUB boot loaders or the Linux kernel, you can remove power and then the microSD card. Your device should now be able to boot from USB ([Section 3.6.4, “Installation from USB media”](#)).

Note that once USB boot mode has been enabled for Raspberry Pi 3 Model B, USB boot mode cannot be disabled again ([Section 3.6.1.2, “OTP memory”](#)).

For more details, refer to the Raspberry Pi Foundation Web site: <https://www.raspberrypi.org/documentation/hardware/raspberrypi/bootmodes/msd.md> 

For the Raspberry Pi Compute Module 3, the setting required is the same, but the deployment of the modified image is a little more complicated.

3.6.2 Lack of a real-time clock

There is no battery-backed Real-Time Clock (RTC) on the Raspberry Pi itself.



Note: Time synchronization

The lack of a Real-Time Clock means that Raspberry Pi devices need to be configured to fetch the time from a network server by Network Time Protocol (NTP).

However, base boards for the Raspberry Pi Compute Modules may feature an RTC.

It is also possible to connect an RTC via the GPIO connector, using Hardware Attached on Top (HATs) or other expansion boards.

Either way, check whether the respective RTC chipset is supported by SUSE Linux Enterprise Server. The connected RTC will need to be described to the operating system via a Device Tree Overlay ([Section 3.6.1.1, "Config.txt"](#)).

Compute Module 4 IO Board

```
dtparam=i2c_vc=on
dtoverlay=i2c-rtc,pcf85063a,i2c_csi_dsi
```

MyPi base board


```
dtparam=i2c1=on
dtoverlay=i2c-rtc,ds1307
```

For other boards and HATs, consult the documentation they are shipped with.

3.6.3 Deploying an appliance image

The most common method to deploy an operating system onto Raspberry Pi hardware is to copy a pre-installed system image onto a boot medium, usually a microSD card. This is the simplest and easiest method.

SUSE provides a preconfigured bootable image of SUSE Linux Enterprise Server for Raspberry Pi hardware. This comes with the Btrfs file system, with compression enabled to improve performance and reduce wear on microSD media.

A microSD card with a minimum size of 8 GB is recommended. Faster cards will give better system performance. On the first boot, the operating system automatically expands the file system to fill the card. This means that the first boot will be substantially slower than subsequent boots. The process of writing the card image onto microSD media is described in the [Raspberry Pi Quick Start \(https://documentation.suse.com/sles/html/SLES-raspberry-pi/article-raspberry-pi.html\)](https://documentation.suse.com/sles/html/SLES-raspberry-pi/article-raspberry-pi.html) .

3.6.4 Installation from USB media

Some models of Raspberry Pi allow booting from USB mass storage devices. This will then allow deploying SUSE Linux Enterprise Server on Raspberry Pi similar to server platforms.

Installation can be performed from a removable USB medium, such as a memory stick, onto a microSD card in the machine's internal slot. Alternatively, it can be performed from a removable USB medium onto another USB medium, such as a USB-connected hard disk.



Note: USB bandwidth limitations

Note that the Ethernet controller on the Raspberry Pi 3 is connected to the device's on-board USB 2.0 bus.

Therefore an operating system running from a disk attached via USB must share the total 480 Mbps bandwidth of the USB 2.0 controller. This will limit performance, and could significantly impact network performance.

This limitation does not apply to the Raspberry Pi 4.

Newer models of Raspberry Pi 3 with BCM2837 B0 silicon (silver instead of black chip), including Raspberry Pi 3 Model B+ and Compute Module 3+, allow booting from USB-connected storage devices by default.

On older models, such as Raspberry Pi 3 Model B or Compute Module 3, USB boot can be enabled by booting from a specially prepared microSD card once. See [Section 3.6.1.2, "OTP memory"](#) for instructions.

3.6.5 Installation from network

Because of the hardware's lack of on-board firmware ([Section 3.6.1, “Boot process”](#)), network-booting the Raspberry Pi using PXE is more complex than with more conventional computers.

The process of setting up a PXE boot server for x86 and Arm is described in the SUSE Best Practices document [How to Set Up a Multi-PXE Installation Server](https://documentation.suse.com/sbp/all/html/SBP-Multi-PXE-Install/index.html) (<https://documentation.suse.com/sbp/all/html/SBP-Multi-PXE-Install/index.html>) [↗](#).

The Raspberry Pi Foundation publishes information on how to boot using PXE one Raspberry Pi from another Raspberry Pi: https://www.raspberrypi.org/documentation/hardware/raspberrypi/bootmodes/net_tutorial.md [↗](#)

3.6.6 More information

For more information, consult the following resources:

SUSE Linux Enterprise Server 15 SP4 Release Notes

For more information about hardware compatibility, supported options and functionality when running on Raspberry Pi hardware, consult the *Boot and Driver Enablement for Raspberry Pi* section of the SUSE Linux Enterprise Server Release Notes:

<https://www.suse.com/releasenotes/aarch64/SUSE-SLES/15-SP4/#aarch64-rpi> [↗](#)

Raspberry Pi Quick Start

<https://documentation.suse.com/sles/15-SP4/html/SLES-raspberry-pi/article-raspberry-pi.html> [↗](#)

openSUSE Hardware Compatibility List: Raspberry Pi 3

The openSUSE project also has information about installing and configuring Raspberry Pi hardware. Much of this also applies to SUSE Linux Enterprise.

See https://en.opensuse.org/HCL:Raspberry_Pi3 [↗](#).

Das U-Boot

More information about Das U-Boot boot loader can be found on the project's GitHub page at <https://github.com/u-boot/u-boot> [↗](#).

4 Installation on IBM POWER

This chapter describes the installation procedure of SUSE Linux Enterprise Server on IBM POWER systems.

4.1 Hardware requirements

To run SUSE Linux Enterprise Server on POWER, your hardware must meet the minimum requirements listed below.

Supported servers

Check the database of SUSE-certified hardware to make sure that your particular hardware configuration is supported. The database is available at <https://www.suse.com/yessearch/Search.jsp>. SUSE Linux Enterprise Server may support additional IBM POWER systems that are not listed. For the latest information, refer to the IBM Information Center for Linux at <https://www.ibm.com/support/knowledgecenter/linuxonibm/liaam/liaamdistros.htm>.

Memory requirements

A minimum of 1024 MB of memory is required for a minimal installation. On machines with more than two processors, add 512 MB per CPU. For remote installations via HTTP or FTP, add another 150 MB. Note that these values are only valid for the installation of the operating system—the actual memory requirement in production depends on the system's workload. For systems running the GNOME desktop environment, a minimum of 2048 MB of memory is required and 4096 MB is recommended.

Hard disk requirements

The disk requirements depend on the type of installation selected and the usage scenario. Normally, a properly working system requires more space than the installation itself. The minimum requirements are as follows.

Installation Scope	Minimum Hard Disk Requirements
Text Mode	1.5 GB
Minimal System	2.5 GB

Installation Scope	Minimum Hard Disk Requirements
GNOME Desktop	3 GB
All patterns	4 GB
Recommended Minimum (no Btrfs snapshots): 10 GB	
Required Minimum (with Btrfs snapshots): 16 GB	
Recommended Minimum (with Btrfs snapshots): 32 GB	

If the root partition is smaller than 10 GB, the installer does not offer a partitioning proposal. In this case, you need to create partitions manually. To avoid this, we recommend to have 10 GB reserved for the root partition. Increase the minimum size to 16 GB if you plan to enable Btrfs snapshots on the root volume (see *Book “Administration Guide”, Chapter 10 “System recovery and snapshot management with Snapper”*).

Before installing SUSE Linux Enterprise Server, make sure that the server has the latest firmware. For the latest firmware, visit IBM FixCentral: <https://www.ibm.com/support/fixcentral/>. Select your system from the Product Group list. Additional software is available from the IBM PowerLinux Tools Repository. For more information on using the IBM PowerLinux Tools Repository, see <https://www.ibm.com/docs/en/linux-on-systems?topic=servers-linux-power-tools-repository>.

4.2 Installing SUSE Linux Enterprise Server for POWER

The following procedure describes how to set up an installation environment. You can skip it if you already have an installation environment ready.

PROCEDURE 4.1: PREPARING AN INSTALLATION ENVIRONMENT

1. Start an SSH session to your HMC and run the `vtmenu` command.
2. Select the desired POWER server and the LPAR. If a serial console session for the chosen LPAR already exists, you need to close it first using the following command:

```
rmvterm -m SERVER -p LPAR
```

3. Reboot the LPAR by creating a new SSH session to the HMC and running the following command:

Note that this command causes a hard reboot of the LPAR. To perform a soft reboot and allow the running tasks to shut down properly, omit the `--immed` flag on the command above.

[illegible]

5. Select Setup Remote IPL (Initial Program Load) by pressing 2 and **Enter**.

6. Select the NIC Adapter for accessing your TFTP server.

7. Select the IP version to be used (for example, IPv4).
8. Select the protocol used to access the TFTP server (for example, 1 for BOOTP).
9. Select IP Parameters by pressing 1 and **Enter** .
10. Configure the required network parameters of the LPAR, including the IP address, the network gateway, and the network mask. In the Server IP Address, specify the IP address of your TFTP server.

```
PowerPC Firmware
Version FW940.01 (VL940_034)
SMS (c) Copyright IBM Corp. 2000,2019 All rights reserved.
-----
IP Parameters
Interpartition Logical LAN: U9008.22L.787FE9A-V8-C2-T1
1. Client IP Address      [10.161.24.65]
2. Server IP Address     [10.161.0.99]
3. Gateway IP Address    [10.161.0.1]
4. Subnet Mask           [255.255.192.0]

-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:|
```

11. Use the **Esc** key to return to the first screen. Select the following entries in the specified order:
 - Select Boot Options
 - Select Install/Boot Device
 - Network
 - BOOTP
12. Select the NIC adapter specified earlier, then choose:
 - Normal Mode Boot
 - Yes
13. When the process starts, you should see a GRUB menu containing a list of images available on the TFTP server.

```
GNU GRUB  version 2.02

+-----+
| linux                                     |
| local                                   |
| ppc64le:SLE-12-SP4-Server-LATEST       |
| *ppc64le:SLE-12-SP5-Server-LATEST      |
| ppc64le:SLE-15-Installer-LATEST        |
| ppc64le:SLE-15-SP1-Installer-LATEST    |
| ppc64le:SLE-15-SP2-Full-LATEST         |
+-----+


Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, `e' to edit the commands
before booting or `c' for a command-line. ESC to return
previous menu.
```

4.3 Installing SUSE Linux Enterprise Server

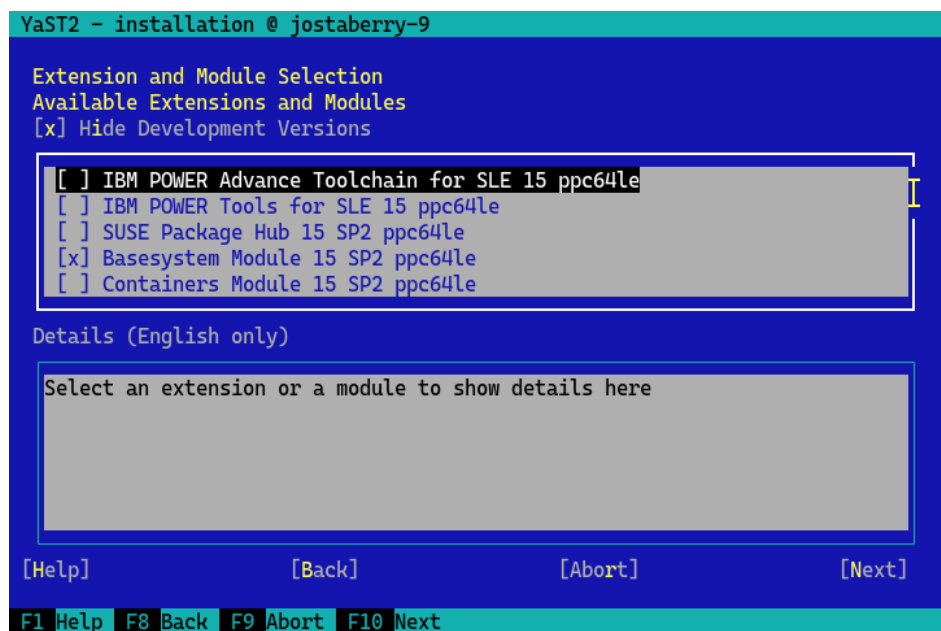
In general, installing SUSE Linux Enterprise Server on POWER is similar to a regular installation procedure.

PROCEDURE 4.2: SUSE LINUX ENTERPRISE SERVER INSTALLATION

1. In the first two steps, you are prompted to choose the desired language and keyboard and to read and agree to the product's license agreement.
2. Next, choose the desired product registration method and complete the registration. If you register the system using the SUSE Customer Center, you are prompted to enable update repositories. Press Yes.



3. To install any modules or extensions, select each one using the arrow keys and pressing **Space** . Depending on what extensions and modules you select, you may be prompted to import GnuPG keys for the associated repositories.



4. Install the desired add-on products. If you choose to install an add-on, you need to specify the installation source for it.

```
YaST2 - installation @ jostaberry-9

Add On Product
[ ] I would like to install an additional Add On Product

( ) Scan Using SLP...
(x) Specify URL...
( ) FTP...
( ) HTTP...
( ) HTTPS...
( ) SMB/CIFS
( ) NFS...
( ) CD...
( ) DVD...
( ) Hard Disk...
( ) USB Mass Storage (USB Stick, Disk)...
( ) Local Directory...
( ) Local ISO Image...

[x] Download repository description files
[Help] [Back] [Abort] [Next]

F1 Help F8 Back F9 Abort F10 Next
```

5. Specify a partition scheme for your installation. To accept the default proposal, press Next or press **Alt** - **N** .

```
YaST2 - installation @ jostaberry-9

Suggested Partitioning [Release Notes ...]

Initial layout proposed with the default Guided Setup settings.
Changes to partitioning:


* Delete partition
  /dev/mapper/3600a098038302d414b2b4d4453474f55-part2 (30.86 GiB)
* Delete partition
  /dev/mapper/3600a098038302d414b2b4d4453474f55-part4 (15.90 GiB)
* Delete partition
  /dev/mapper/3600a098038302d414b2b4d4453474f55-part3 (17.24 GiB)
* Create partition
  /dev/mapper/3600a098038302d414b2b4d4453474f55-part2 (30.86 GiB)

[Guided Setup]
[Expert Partitioner]

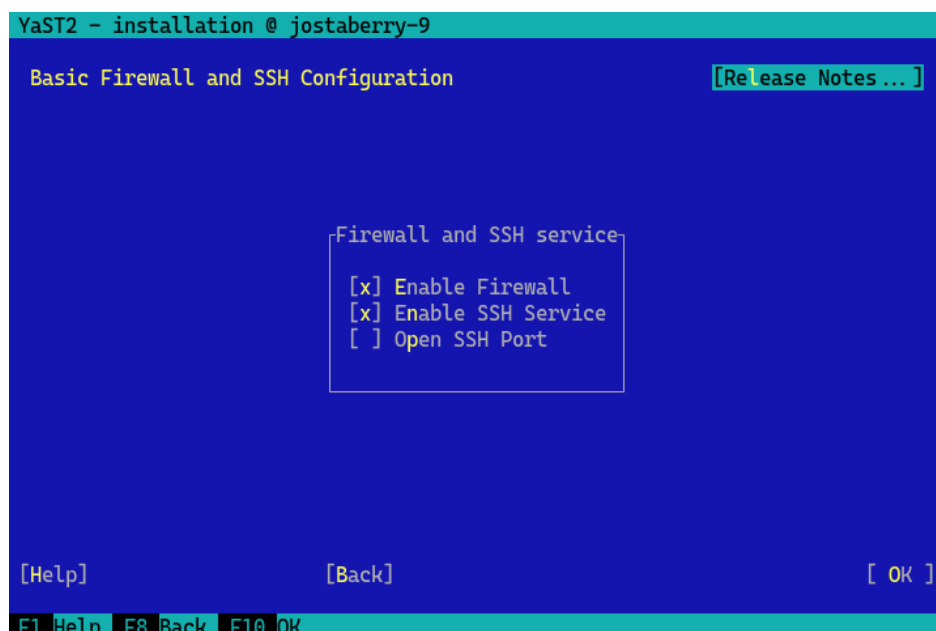
[Help] [Back] [Abort] [Next]

F1 Help F8 Back F9 Abort F10 Next
```

6. Choose the system role suitable for your particular scenario.






- The next few screens allow you to specify the appropriate time zone, and create a user. If you choose not to create a user, you are prompted to specify a root password.
- In the installation summary screen, make sure the SSH service is enabled and open an SSH port. To do this, press Change, go to the Basic Firewall and SSH Configuration screen, and enable the appropriate options. Press OK.



- Confirm the installation configuration, and press Install to start the installation process.

4.4 More information

Further information on IBM PowerLinux is available from SUSE and IBM:

- The SUSE Support Knowledge Base at <https://www.suse.com/support/kb/>  is a help tool for assisting customers in solving problems. Search the knowledge base on SUSE Linux Enterprise Server using relevant search terms.
- Find security alerts at <https://www.suse.com/support/security/> . SUSE also maintains two security-related mailing lists:
 - suse-security — General discussion of security topics related to Linux and SUSE. All security alerts for SUSE Linux Enterprise Server are sent to this list.
 - suse-security-announce — The SUSE mailing list exclusively for security alerts.
- To participate in the linuxppc-dev mailing list, register using the forms at <https://lists.ozlabs.org/listinfo/linuxppc-dev/> .

5 Installation on IBM Z and LinuxONE

This chapter describes the procedure for preparing the installation of SUSE® Linux Enterprise Server on IBM Z. It provides all information needed to prepare the installation on the LPAR and z/VM side.

5.1 System requirements

This section provides basic information about the system requirements, MicroCode level, and software for IBM Z.

5.1.1 Hardware

SUSE Linux Enterprise Server runs on the following platforms:

- IBM zEnterprise EC12 (zEC12) (2827)
- IBM zEnterprise BC12 (zBC12) (2828)
- IBM z Systems z13 (2964)
- IBM z Systems z13s (2965)
- IBM z Systems z14 (3906)
- IBM z Systems z14 ZR1 (3907)
- IBM z Systems z15 T01 (8561)
- IBM z Systems z15 T02 (8562)
- IBM z Systems z16 A01 (3931)
- IBM LinuxONE Emperor (2964)
- IBM LinuxONE Rockhopper (2965)
- IBM LinuxONE Emperor II (3906)
- IBM LinuxONE Rockhopper II (3907)
- IBM LinuxONE III LT1 (8561)

- IBM LinuxONE III LT2 (8562)
- IBM LinuxONE Emperor 4 (3931)

5.1.1.1 Memory requirements

Different installation methods have different memory requirements during installation. At least 1 GB of memory is recommended for the text-mode installation under z/VM, LPAR, and KVM. Installation in the graphical mode requires at least 1.5 GB of memory.



Note: Memory requirements with remote installation sources

A minimum of 512 MB of memory is required for installation from NFS, FTP, and SMB installation sources, or when VNC is used. Keep in mind that memory requirements also depend on the number of devices visible to the z/VM guest or the LPAR image. Installation with many accessible devices (even if unused for the installation) may require more memory.

5.1.1.2 Disk space requirements

The disk requirements depend largely on the installation. To have a properly functioning system, you normally need more space than required by the installation software. Minimal requirements for the available installation types are as follows:

Installation Type	Minimum Hard Disk Requirements
Text Mode	1.5 GB
Minimal System	2.5 GB
GNOME Desktop	3 GB
All patterns	4 GB
Recommended Minimum (no Btrfs snapshots): 10 GB	
Required Minimum (with Btrfs snapshots): 16 GB	
Recommended Minimum (with Btrfs snapshots): 32 GB	

5.1.1.3 Network connection

A network connection is needed to communicate with your SUSE Linux Enterprise Server system. This can be one or several of the following connections or network cards:

- OSA Express Ethernet (including Fast and Gigabit Ethernet)
- HiperSockets or Guest LAN
- 10 GBE, VSWITCH
- RoCE (RDMA over Converged Ethernet)


The following interfaces are still included, but no longer supported:

- CTC (or virtual CTC)
- ESCON
- IP network interface for IUCV

For installations under KVM, make sure the following requirements are met to enable the VM Guest to access the network transparently:

- The virtual network interface is connected to a host network interface.
- The host network interface is connected to a network that the virtual server will join.
- If the host is configured to have a redundant network connection by grouping two independent OSA network ports into a bonded network interface, the identifier for the bonded network interface is `bond0`. If more than one bonded interface exists, it is `bond1`, `bond2`, etc.
- A non-redundant network connection setup requires the identifier of the single network interface. The identifier has the following format: `enccw0.0.NNNN`, where `NNNN` is the device number of the desired network interface.

5.1.2 MicroCode Level, APARs, and fixes

Documentation about restrictions and requirements for this release of SUSE Linux Enterprise Server be found on IBM developerWorks at <https://developer.ibm.com/technologies/linux/> . We recommend to use the highest service level available. Contact IBM support for minimum requirements.

For z/VM, the following versions are supported:

- z/VM 6.4
- z/VM 7.1
- z/VM 7.2
- z/VM 7.3

Since it might be necessary to activate the VM APARs before installing the new MicroCode levels, clarify the order of installation with IBM support.

5.1.3 Software


When installing SUSE Linux Enterprise Server via non-Linux-based NFS or FTP, you might experience problems with NFS or FTP server software. The Windows* standard FTP server can cause errors, so we recommend performing installation via SMB on these machines.

To connect to the SUSE Linux Enterprise Server installation system, one of the following methods is required (SSH or VNC are recommended):

SSH with terminal emulation (xterm compatible)

SSH is a standard Unix tool that is present on most Unix or Linux systems. For Windows, you can use the Putty SSH client.

VNC client

For Linux, the `vncviewer` VNC client is included in SUSE Linux Enterprise Server as part of the `tightvnc` package. For Windows, TightVNC is also available. Download it from <https://www.tightvnc.com/> .

X server

Find a suitable X server implementation on any Linux or Unix workstation. There are many commercial X Window System environments for Windows and macOS*. Some can be downloaded as free trial versions.



Tip: More information

Before installing SUSE Linux Enterprise Server on IBM Z, consult the `README` file located in the root directory of the first installation medium of SUSE Linux Enterprise Server. The file complements this documentation.

5.2 General information

This section covers the different installation types and how to do an IPL for the first installation.

5.2.1 Installation types

This section gives an overview of the different types of installation possible with SUSE Linux Enterprise Server for IBM Z. SUSE Linux Enterprise Server can be installed in an *LPAR*, as a guest within *z/VM*, or as a guest within *KVM*.

Depending on the mode of installation (LPAR or *z/VM*), there are different possibilities for starting the installation process and IPLing the installed system.

5.2.1.1 LPAR

If you install SUSE Linux Enterprise Server for IBM Z into a logical partition (LPAR), assign memory and processors to the instance. Installing into LPAR is recommended for highly loaded production machines. Running in LPAR also makes higher security standards available. Networking between LPARs is possible over external interfaces or HiperSockets. In case you plan to use your installation for virtualization with KVM, installing into LPAR is highly recommended.

5.2.1.2 z/VM

Running SUSE Linux Enterprise Server for IBM Z in *z/VM* means that SUSE Linux Enterprise Server is a guest system within *z/VM*. An advantage of this mode is that you have full control over SUSE Linux Enterprise Server from *z/VM*. This is very helpful for kernel development or kernel-based debugging. It is also very easy to add or remove hardware to and from Linux guests. Creating additional SUSE Linux Enterprise Server guests is simple and you can run hundreds of Linux instances simultaneously.

5.2.1.3 KVM guest

Being able to install SUSE Linux Enterprise Server for IBM Z as a KVM guest requires a KVM host server instance installed into LPAR. For details on the guest installation, refer to [Procedure 5.3, “Overview of a KVM guest installation”](#).

5.2.2 IPL options

This section provides the information needed to do an IPL for the first installation. Depending on the type of installation, different options need to be used. The VM reader, load from CD-ROM or server and load from an SCSI-attached DVD-ROM options are discussed. Installing the software packages, which is done over the network, does not require the IPL medium.

5.2.2.1 VM reader

To IPL from a VM reader, transfer the necessary files into the reader first. For convenience of administration, it is recommended to create a user `linuxmnt` that owns a minidisk with the files and scripts needed for IPL. This minidisk is then accessed read-only by the Linux guests. For details, see [Section 5.3.4.2.1, “IPL from the z/VM reader”](#).

5.2.2.2 Load from removable media or server

For IPLing into an LPAR, load the kernel image directly from the SE's or the HMC's CD/DVD-ROM device or from any remote system accessible through FTP. This function can be performed from the HMC. The installation process requires a file with a mapping of the location of the installation data in the file system and the memory locations to which to copy the data.

For SUSE Linux Enterprise Server, there are two such files. Both are located in the root directory of the first installation medium:

- `suse.ins`, for which to work you need to set up network access in `Linuxrc` before starting the installation.
- `susehmc.ins` which allows installing without network access.

In the left navigation pane of the HMC expand *Systems Management* > *Systems* and select the mainframe system you want to work with. Choose the LPAR where you want to boot SUSE Linux Enterprise Server from the table of LPARs and select *Load from Removable Media or Server*.

Now either choose *Hardware Management Console CD-ROM/DVD* or *FTP Source*. If having chosen the latter option, provide the servers address or name and your credentials. If the appropriate `.ins` file is not located in the root directory of the server, provide the path to this file. Continue to the *Select the software to load* menu and select the appropriate `.ins` entry. Start the installation with *OK*.

5.2.2.3 Load from SCSI-attached DVD

To IPL from an SCSI DVD, you need access to an FCP adapter connected to a DVD drive. You need the values for WWPN and LUN from the SCSI drive. For details, see [Section 5.3.4.1.2, “IPL from FCP-attached SCSI DVD”](#).

5.2.2.4 Load from the network with zPXE

IPLing from the Network with zPXE requires a Cobbler server providing the kernel, RAM disk and a parmfile. It is initiated by running the ZPXE EXEC script. See [Section 5.3.1.3, “Using a Cobbler server for zPXE”](#) for details. zPXE is only available on z/VM.

5.3 Preparing for installation

This chapter explains how to make the data accessible for installation, install SUSE Linux Enterprise Server using different methods, and prepare and use the IPL of the SUSE Linux Enterprise Server installation system. The chapter also provides information about network configuration and network installation.

5.3.1 Making the installation data available

This section provides detailed information about making the SUSE Linux Enterprise Server IBM Z installation data accessible for installation. Depending on your computer and system environment, choose between NFS or FTP installation. If you are running Microsoft Windows workstations in your environment, you can use the Windows network (including the SMB protocol) to install SUSE Linux Enterprise Server on your IBM Z system.



Tip: IPL from DVD

It is possible to IPL from DVD and use the DVD as the installation medium. This is very convenient if you have restrictions setting up an installation server providing installation media over your network. The prerequisite is an FCP-attached SCSI DVD Drive.



Note: No installation from hard disk

It is not possible to perform installation from a hard disk by putting the content of the DVD to a partition on a DASD.

5.3.1.1 Using a Linux workstation or SUSE Linux Enterprise Server DVD

You can use a Linux workstation in your computer environment to provide the installation data to the IBM Z installation process by NFS or FTP. If the Linux workstation runs SUSE Linux Enterprise Server, you can set up an installation server (NFS or FTP) using the YaST *Installation Server* module as described in [Section 17.1, “Setting up an installation server using YaST”](#).



Important: Exporting mounted devices with NFS

Exporting the file system root (/) does not automatically export the mounted devices, such as DVD. Therefore, you need to explicitly name the mount point in /etc/exports:

```
/media/dvd *(ro)
```

After changing this file, restart the NFS server with the command **`sudo systemctl restart nfsserver`**.

Setting up an FTP server on a Linux system involves the installation and configuration of server software like `vsftpd`. If you are using SUSE Linux Enterprise Server, refer to *Book “Administration Guide”, Chapter 43 “Setting up an FTP server with YaST”* for installation instructions. Downloading the installation data via anonymous login is not supported, therefore you need to configure the FTP server to support user authentication.

5.3.1.1.1 SUSE Linux Enterprise Server on DVD

The first installation medium of the SUSE Linux Enterprise Server for IBM Z contains a bootable Linux image for Intel-based workstations and an image for IBM Z.

For Intel-based workstations, boot from this medium. When prompted, choose the desired answer language and keyboard layout and select *Start rescue system*. You need at least 64 MB RAM for this. No disk space is needed, because the entire rescue system resides in the workstation's RAM. This approach requires setting up the networking of the workstation manually.

For IBM Z, IPL your LPAR/VM guest from this medium as described in [Section 5.3.4.1.2, “IPL from FCP-attached SCSI DVD”](#). After entering your network parameters, the installation system treats the medium as the source of installation data. Because IBM Z cannot have an X11-capable terminal attached directly, choose between VNC or SSH installation. Refer to [Section 12.3, “Monitoring installation via VNC”](#) or [Section 12.4, “Monitoring installation via SSH”](#) for more information. SSH also provides a graphical installation by tunneling the X connection through SSH with **ssh -X**.



Important: **ssh -X** connections between different architectures

By default, recent versions of the X.org and Xwayland servers do not accept connections from clients on different architectures. If you connect to the IBM Z machine from a AMD64/Intel 64 workstation with **ssh -X**, you will likely see the error message: “Prohibited client endianness, see the Xserver man page”.

To enable X connections between different architectures, create the file `/etc/X11/xorg.conf.d/99-byte-swapping.conf` with the following content:

```
Section "ServerFlags"
    Option "AllowByteSwappedClients" "on"
EndSection
```

Restart your X.org or Xwayland server to apply the configuration change:

```
> sudo systemctl restart display-manager.service
```

5.3.1.2 Using a Microsoft Windows workstation

You can use a Microsoft Windows workstation on your network to make the installation media available. The easiest way to do this is to use the SMB protocol. Make sure to activate *SMB over TCP/IP* as this enables the encapsulation of SMB packages into TCP/IP packages. Find details in the Windows online help or other Windows-related documentation that covers networking.

5.3.1.2.1 Using SMB

To make the installation media available with SMB, insert the USB flash drive with SLE-15-SP6-Online-*ARCH*-GM-media1.iso into the USB port of the Windows workstation. Then create a new share using the USB flash drive's letter and make it available for everyone in the network.

The installation path in YaST can be:

```
smb://DOMAIN;USER:PW@SERVERNAME/SHAREPATH
```

Where the placeholders mean:

DOMAIN

Optional workgroup or active directory domain.

USER,

PW

Optional user name and password of a user who can access this server and its share.

SERVERNAME

The name of the server that hosts the share(s).

SHAREPATH

The path to the share(s).

5.3.1.2.2 With NFS

Refer to the documentation provided with the third party product that enables NFS server services for your Windows workstation. The USB flash drive containing the SLE-15-SP6-Online-ARCH-GM-media1.iso medium must be in the available NFS path.

5.3.1.2.3 Using FTP

Refer to the documentation provided with the third-party product that is enabling FTP server services on your Windows workstation. The USB flash drive containing the SLE-15-SP6-Online-ARCH-GM-media1.iso medium must be in the available FTP path.

The FTP server that is bundled with certain Microsoft Windows releases implements only a subset of the FTP commands, and it is not suitable for providing the installation data. In this case, use a third-party FTP server that offers the required functionality.

5.3.1.2.4 Using an FCP-attached SCSI DVD drive

After you IPLed from the SCSI DVD as described in [Section 5.3.4.1.2, “IPL from FCP-attached SCSI DVD”](#), the installation system uses the DVD as the installation medium. In this case, you do not need the installation media on an FTP, NFS, or SMB server. However, you need the network configuration data for your SUSE Linux Enterprise Server, because you must set up the network during the installation to perform a graphical installation via VNC or by X.

5.3.1.3 Using a Cobbler server for zPXE

IPLing from the network requires a Cobbler server to provide the kernel, initrd, and the installation data. Preparing the Cobbler server requires the following steps:

- [Section 5.3.1.3.1](#)
- [Section 5.3.1.3.2](#)
- [Section 5.3.1.3.3](#)
- [Section 5.3.1.3.4](#)

5.3.1.3.1 Importing the installation data

Importing the media requires the installation source to be available on the Cobbler server—either from USB flash drive or from a network source. Run the following command to import the data:

```
> sudo cobbler import --path=PATH ❶ --name=IDENTIFIER ❷ --arch=s390x
```

- ❶ Mount point of the installation data.
- ❷ A string identifying the imported product, for example “sles15_s390x”. This string is used as the name for the subdirectory where the installation data is copied to. On a Cobbler server running on SUSE Linux Enterprise this is /srv/www/cobbler/ks_mirror/IDENTIFIER. This path may be different if Cobbler runs on another operating system.

5.3.1.3.2 Adding a distribution

Adding a distribution allows Cobbler to provide the kernel and the initrd required to IPL via zPXE. Run the following command on the Cobbler server to add SUSE Linux Enterprise Server for IBM Z:

```
> sudo cobbler distro add --arch=s390 --breed=suse --name="IDENTIFIER" ❶ \  
  --os-version=sles15 ❷ \  
  --initrd=/srv/www/cobbler/ks_mirror/IDENTIFIER/boot/s390x/initrd ❸ \  
  --kernel=/srv/www/cobbler/ks_mirror/IDENTIFIER/boot/s390x/linux ❹ \  
  --kopts="install=http://cobbler.example.com/cobbler/ks_mirror/IDENTIFIER" ❺
```

- ❶ Unique identifier for the distribution, for example “SLES 15 SP6 IBM Z”.
- ❷ Operating system identifier. Use `sles15`.
- ❸ Path to the initrd. The first part of the path (`/srv/www/cobbler/ks_mirror/IDENTIFIER/`) depends on the location of the imported data and the subdirectory name you chose when importing the installation data.
- ❹ Path to the kernel. The first part of the path (`/srv/www/cobbler/ks_mirror/IDENTIFIER/`) depends on the location of the imported data and the subdirectory name you chose when importing the installation data.
- ❺ URL to the installation directory on the Cobbler server.

5.3.1.3.3 Adjusting the profile

Adding a distribution (see [Section 5.3.1.3.2, “Adding a distribution”](#)) automatically generates a profile with the corresponding `IDENTIFIER`. Use the following command to make a few required adjustments:

```
> sudo cobbler distro edit \  
  --name=IDENTIFIER ❶ --os-version=sles10 ❷ --ksmeta="" ❸ \  
  --kopts="install=http://cobbler.example.com/cobbler/ks_mirror/IDENTIFIER" ❹
```

- ❶ Identifier for the profile. Use the string specified when added the distribution.
- ❷ Operating system version. Distribution to which the profile should apply. Use the string specified with `--name=IDENTIFIER` in the importing step.
- ❸ Option required for templating Kickstart files. Since it is not used for SUSE, leave it empty.
- ❹ Space-separated list of kernel parameters. It must include at least the `install` parameter.

5.3.1.3.4 Adding systems

The last step is to add systems to the Cobbler server. This step must be performed for every IBM Z guest that should boot via zPXE. Guests are identified by their z/VM user ID (in the following example, the ID “linux01”). Note that the ID must be lowercase. To add a system, run the following command:

```
> sudo cobbler system add --name=linux01 --hostname=linux01.example.com \  
--profile=IDENTIFIER --interface=qdio \  
--ip-address=192.168.2.103 --subnet=192.168.2.255 --netmask=255.255.255.0 \  
--name-servers=192.168.1.116 --name-servers-search=example.com \  
--gateway=192.168.2.1 --kopts="KERNEL_OPTIONS"
```

The `--kopts` option allows you to specify the kernel and installation parameters that are usually specified in the parmfile. Specify the parameters using the following format: `PARAMETER1=VALUE1` `PARAMETER2=VALUE2`. The installer prompts for missing parameters. For a fully-automated installation, you need to specify all parameters for networking, DASDs and provide an AutoYaST file. Below is an example for a guest equipped with an OSA interface using the same network parameters as above.

```
--kopts=" \  
AutoYaST=http://192.168.0.5/autoinst.xml \  
Hostname=linux01.example.com \  
Domain=example.com \  
HostIP=192.168.2.103 \  
Gateway=192.168.2.1 \  
Nameserver=192.168.1.116 \  
Searchdns=example.com \  
InstNetDev=osa; \  
Netmask=255.255.255.0 \  
Broadcast=192.168.2.255 \  
OsaInterface=qdio \  
Layer2=0 \  
PortNo=0 \  
ReadChannel=0.0.0700 \  
WriteChannel=0.0.0701 \  
DataChannel=0.0.0702 \  
DASD=600"
```

5.3.1.4 Installing from a USB Flash Drive of the HMC

Installation of SUSE Linux Enterprise Server on IBM Z servers usually requires a network installation source. If this requirement cannot be fulfilled, SUSE Linux Enterprise Server allows you to use the USB flash drive of the Hardware Management Console (HMC) as an installation source for installation on an LPAR.

To perform installation from the USB flash drive of the HMC, proceed as follows:

- Add `install=hmc:/` to the `parmfile` (see [Section 5.5, “The `parmfile`—automating the system configuration](#)”) or kernel options.
- In the manual-mode installation using `linuxrc`, choose *Start Installation*, then *Installation*, and then *Hardware Management Console*. The installation medium must be in the HMC.



Important: Configure network

Before starting the installation, specify a network configuration in `linuxrc`. You cannot do this via boot parameters, and it is very likely that you will need network access. In `linuxrc`, go to *Start Installation*, then choose *Network Setup*.



Important: Linux system must boot first

Before granting access to the media in the USB flash drive of the HMC, wait until the Linux system is booted. IPLing can disrupt the connection between the HMC and the LPAR. If the first attempt to use the described method fails, you can grant the access and retry the option `HMC`.



Note: Installation repository

The USB flash drive is not kept as an installation repository, as the installation is a one-time procedure. If you need an installation repository, register and use the online repository.

5.3.2 Installation types

This section describes SUSE Linux Enterprise Server installation steps for each installation mode. When the preparation steps described in the previous chapters have been completed, follow the overview of the desired installation mode.

As described in [Section 5.3.1, “Making the installation data available”](#), there are three different installation modes for Linux on IBM Z: LPAR, z/VM, and KVM guest installation.

PROCEDURE 5.1: OVERVIEW OF AN LPAR INSTALLATION

1. Prepare the devices needed for installation. See [Section 5.3.3.1, “Preparing the IPL of an LPAR installation”](#).
2. IPL the installation system. See [Section 5.3.4.1, “IPLing an LPAR installation”](#).
3. Configure the network. See [Section 5.3.5, “Network configuration”](#).
4. Connect to the SUSE Linux Enterprise Server installation system. See [Section 5.3.6, “Connecting to the SUSE Linux Enterprise Server installation system”](#).
5. Start the installation using YaST and IPL the installed system. See [Chapter 9, Installation steps](#).

PROCEDURE 5.2: INSTALLATION OVERVIEW OF Z/VM INSTALLATION

1. Prepare the devices needed for installation. See [Section 5.3.3.2.1, “Adding a Linux guest using dirMaint”](#).
2. IPL the installation system. See [Section 5.3.4.2, “IPLing a z/VM installation”](#).
3. Configure the network. See [Section 5.3.5, “Network configuration”](#).
4. Connect to the SUSE Linux Enterprise Server installation system. See [Section 5.3.6, “Connecting to the SUSE Linux Enterprise Server installation system”](#).
5. Start the installation using YaST and IPL the installed system. See [Chapter 9, Installation steps](#).

PROCEDURE 5.3: OVERVIEW OF A KVM GUEST INSTALLATION

1. Create a virtual disk image and write a domain XML file. See [Section 5.3.3.3, “Preparing the IPL of a KVM guest installation”](#).

2. Prepare the installation target and IPL the VM Guest. See [Section 5.3.4.3, “IPLing a KVM guest installation”](#).
3. [Section 5.3.5.3, “Set up the network and select the installation source”](#).
4. Connect to the SUSE Linux Enterprise Server installation system. See [Section 5.3.6, “Connecting to the SUSE Linux Enterprise Server installation system”](#).
5. Start the installation using YaST and IPL the installed system. See [Chapter 9, Installation steps](#).

5.3.3 Preparing the IPL of the SUSE Linux Enterprise Server installation system

5.3.3.1 Preparing the IPL of an LPAR installation

Configure your IBM Z system to start in ESA/S390 or Linux-only mode with an appropriate activation profile and IOCDS. For further information, refer to the IBM documentation. Continue as described in [Section 5.3.4.1, “IPLing an LPAR installation”](#).

5.3.3.2 Preparing the IPL of a z/VM installation

5.3.3.2.1 Adding a Linux guest using dirMaint

The first step is to attach and format one or multiple DASDs in the system to be used by the Linux guest in z/VM. Next, create a new user in z/VM. The example shows the directory for a user `LINUX1` with the password `LINPWD`, 1 GB of memory (extendable up to 2 GB), several minidisks (MDISK), two CPUs, and an OSA QDIO device.



Tip: Assigning memory to z/VM guests

When assigning memory to a z/VM guest, make sure that the memory size is adequate for the preferred installation type, as described in [Section 5.1.1.1, “Memory requirements”](#). To set the memory size to 1 GB, use the command **CP DEFINE STORAGE 1G**. After the installation has finished, reset the memory size to the desired value.

EXAMPLE 5.1: CONFIGURATION OF A Z/VM DIRECTORY

```
USER LINUX1 LINPWD 1024M 2048M G
*
* _____
* LINUX1
* _____
* This VM Linux guest has two CPUs defined.

CPU 01 CPUID 111111
CPU 02 CPUID 111222
IPL CMS PARM AUTOCR
IUCV ANY
IUCV ALLOW
MACH ESA 10
OPTION MAINTCCW RMCHINFO
SHARE RELATIVE 2000
CONSOLE 01C0 3270 A
SPOOL 000C 2540 READER *
SPOOL 000D 2540 PUNCH A
SPOOL 000E 3203 A
* OSA QDIO DEVICE DEFINITIONS
DEDICATE 9A0 9A0
DEDICATE 9A1 9A1
DEDICATE 9A2 9A2
*
LINK MAINT 0190 0190 RR
LINK MAINT 019E 019E RR
LINK MAINT 019D 019D RR
* MINIDISK DEFINITIONS
MDISK 201 3390 0001 0050 DASD40 MR ONE4ME TW04ME THR4ME
MDISK 150 3390 0052 0200 DASD40 MR ONE4ME TW04ME THR4ME
MDISK 151 3390 0253 2800 DASD40 MR ONE4ME TW04ME THR4ME
```

This example uses minidisk 201 as the guest's home disk. Minidisk 150 with 200 cylinders is the Linux swap device. Disk 151 with 2800 cylinders holds the Linux installation.

As user MAINT, add the guest to the user directory with **DIRM FOR LINUX1 ADD**. Enter the name of the guest (LINUX1) and press **F5**. Set up the environment of the user with:

```
DIRM DIRECT
DIRM USER WITHPASS
```

The last command returns a reader file number. This number is needed for the next command:

```
RECEIVE <number> USER DIRECT A (REPL)
```

You can now log in on the guest as user LINUX1.

If you do not have the `dirmaint` option available, refer to the IBM documentation on how to set up this user.

Proceed with [Section 5.3.4.2, “IPLing a z/VM installation”](#).

5.3.3.3 Preparing the IPL of a KVM guest installation

A KVM guest installation requires a domain XML file that specifies the virtual machine and at least one virtual disk image for the installation.

5.3.3.3.1 Create a virtual disk image

By default, libvirt searches for disk images in `/var/lib/libvirt/images/` on the VM Host Server. Although images can also be stored anywhere on the file system, it is recommended to store all images in a single location for easier maintainability. To create an image, log in to the KVM host server and run the following command:

```
qemu-img create -f qcow2 /var/lib/libvirt/images/s12lin_qcow2.img 10G
```

This creates a qcow2 image with a size of 10 GB in `/var/lib/libvirt/images/`. For more information, refer to *Book “Virtualization Guide”, Chapter 36 “Guest installation”, Section 36.2 “Managing disk images with `qemu-img`”*.

5.3.3.3.2 Write a domain XML file

A domain XML file is used to define the VM Guest. To create the domain XML file, open an empty file `s15-1.xml` with an editor and create a file like in the following example.

EXAMPLE 5.2: EXAMPLE DOMAIN XML FILE

The following example creates a VM Guest with a single CPU, 1 GB RAM, and the virtual disk image created in the previous section ([Section 5.3.3.3.1, “Create a virtual disk image”](#)). It assumes that the virtual server is attached to the host network interface `bond0`. Change the source devices element to match your network setup.

```
<domain type="kvm">
  <name>s15-1</name>
  <description>Guest-System SUSE SLES15</description>
  <memory>1048576</memory>
  <vcpu>1</vcpu>
```

```

<os>
  <type arch="s390x" machine="s390-ccw-virtio">hvm</type>
  <!-- Boot kernel - remove 3 lines after successful installation -->
  <kernel>/var/lib/libvirt/images/s15-kernel.boot</kernel>
  <initrd>/var/lib/libvirt/images/s15-initrd.boot</initrd>
  <cmdline>linuxrcstderr=/dev/console</cmdline>
</os>
<iothreads>1</iothreads>
<on_poweroff>destroy</on_poweroff>
<on_reboot>restart</on_reboot>
<on_crash>preserve</on_crash>
<devices>
  <emulator>/usr/bin/qemu-system-s390x</emulator>
  <disk type="file" device="disk">
    <driver name="qemu" type="qcow2" cache="none" iothread="1" io="native"/>
    <source file="/var/lib/libvirt/images/s15lin_qcow2.img"/>
    <target dev="vda" bus="virtio"/>
  </disk>
  <interface type="direct">
    <source dev="bond0" mode="bridge"/>
    <model type="virtio"/>
  </interface>
  <console type="pty">
    <target type="sclp"/>
  </console>
</devices>
</domain>

```

5.3.4 IPLing the SUSE Linux Enterprise Server installation system

5.3.4.1 IPLing an LPAR installation

There are different ways to IPL SUSE Linux Enterprise Server into an LPAR. The preferred way is to use the *Load from CD-ROM or server* feature of the SE or HMC.

5.3.4.1.1 IPL from DVD-ROM

Mark the LPAR to install and select *Load from CD-ROM or server*. Leave the field for the file location blank or enter the path to the root directory of the first DVD-ROM and select *Continue*. Keep the default selection in the list of options that appears. *Operating system messages* should now show the kernel boot messages.

5.3.4.1.2 IPL from FCP-attached SCSI DVD

You can use the *Load* procedure by selecting *SCSI* as *Load type* to IPL from SCSI. Enter the WWPN (Worldwide port name) and LUN (Logical unit number) provided by your SCSI bridge or storage (16 digits—do not omit the trailing 0s). The boot program selector must be 2. Use your FCP adapter as *Load address* and perform an IPL.

5.3.4.2 IPLing a z/VM installation

This section describes IPLing the installation system to install SUSE Linux Enterprise Server for IBM Z on a z/VM system.

5.3.4.2.1 IPL from the z/VM reader

You need a working TCP/IP connection and an FTP client program within your newly-defined z/VM guest to transfer the installation system via FTP. Setting up TCP/IP for z/VM is beyond the scope of this manual. Refer to the appropriate IBM documentation.

Log in as the z/VM Linux guest to IPL. Make the content of the directory `/boot/s390x` of the Unified Installer (medium 1) available via FTP within your network. From this directory, get the files `linux`, `initrd`, `parmfile`, and `sles.exec`. Transfer the files with a fixed block size of 80 characters. Specify it with the FTP command **locsite fix 80**. It is important to copy `linux` (the Linux kernel) and `initrd` (the installation image) as binary files, so use the `binary` transfer mode. `parmfile` and `sles.exec` need to be transferred in ASCII mode.

The following example shows the required steps. This particular scenario assumes that the required files are accessible from an FTP server at the IP address `192.168.0.3` and the login is `lininst`.

EXAMPLE 5.3: TRANSFERRING THE BINARIES VIA FTP

```
FTP 192.168.0.3
VM TCP/IP FTP Level 530
Connecting to 192.168.0.3, port 21
220 ftpserver FTP server (Version wu-2.4.2-academ[BETA-18])(1)
Thu Feb 11 16:09:02 GMT 2010) ready.
USER
lininst
331 Password required for lininst
PASS
*****
230 User lininst logged in.
```



```

Command:
binary
200 Type set to I
Command:
locsite fix 80
Command:
get /media/dvd1/boot/s390x/linux sles.linux
200 PORT Command successful
150 Opening BINARY mode data connection for /media/dvd1/boot/s390x/linux
(10664192 bytes)
226 Transfer complete.
10664192 bytes transferred in 13.91 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
get /media/dvd1/boot/s390x/initrd sles.initrd
200 PORT Command successful
150 Opening BINARY mode data connection for /media/dvd1/boot/s390x/initrd
(21403276 bytes)
226 Transfer complete.
21403276 bytes transferred in 27.916 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
ascii
200 Type set to A
Command:
get /media/dvd1/boot/s390x/parmfile sles.parmfile
150 Opening ASCII mode data connection for /media/dvd1/boot/s390x/parmfile
(5 bytes)
226 Transfer complete.
5 bytes transferred in 0.092 seconds.
Transfer rate 0.05 Kbytes/sec.
Command:
get /media/dvd1/boot/s390x/sles.exec sles.exec
150 Opening ASCII mode data connection for /media/dvd1/boot/s390x/sles.exec
(891 bytes)
226 Transfer complete.
891 bytes transferred in 0.097 seconds.
Transfer rate 0.89 Kbytes/sec.
Command:
quit

```

Use the REXX script `sles.exec` you downloaded to IPL the Linux installation system. This script loads the kernel, `parmfile`, and the initial RAM disk into the reader for IPL.

EXAMPLE 5.4: SLES.EXEC

```

/* REXX LOAD EXEC FOR SUSE LINUX S/390 VM GUESTS      */

```

```

/* LOADS SUSE LINUX S/390 FILES INTO READER          */
SAY ''
SAY 'LOADING SLES FILES INTO READER...'
'CP CLOSE RDR'
'PURGE RDR ALL'
'SPOOL PUNCH * RDR'
'PUNCH SLES LINUX A (NOH'
'PUNCH SLES PARMFILE A (NOH'
'PUNCH SLES INITRD A (NOH'
'IPL 00C'

```

Using the script, you can IPL the SUSE Linux Enterprise Server installation system with the command **sles**. The Linux kernel then starts and outputs its boot messages.

To continue the installation, proceed to [Section 5.3.5, “Network configuration”](#).

5.3.4.2.2 IPL from FCP-attached SCSI DVD

To IPL in z/VM, prepare the SCSI IPL process by using the SET LOADDEV parameter:

```
SET LOADDEV PORTNAME 200400E8 00D74E00 LUN 00020000 00000000 BOOT 2
```

After setting the LOADDEV parameter with the appropriate values, IPL your FCP adapter, for example:

```
IPL FC00
```

To continue the installation, proceed with [Section 5.3.5, “Network configuration”](#).

5.3.4.2.3 IPL from a Cobbler server with zPXE

To IPL from a Cobbler server with zPXE, you need to transfer the `zpxe.rexx` script via FTP from the Cobbler server to your z/VM guest. To do this, the z/VM guest needs a working TCP/IP connection and an FTP client program.

Log in as the z/VM Linux guest to IPL and transfer the script with a fixed size of 80 characters in ASCII mode (see [Example 5.3, “Transferring the binaries via FTP”](#) for an example). The `zpxe.rexx` script is available on the Unified Installer DVD at `/boot/s390x/zpxe.rexx` or on a SLE Cobbler server at `/usr/share/doc/packages/s390-tools/zpxe.rexx`.

`zpxe.rexx` is supposed to replace the `PROFILE EXEC` of your guest. Make a backup copy of the existing `PROFILE EXEC` and rename `ZPXE REXX` to `PROFILE EXEC`. Alternatively, call `ZPXE REXX` from the existing `PROFILE EXEC` by adding the `'ZPXE REXX'` line to it.

The last step is to create a configuration file `ZPXE CONF` that instructs `ZPXE REXX` which Cobbler server to contact and which disk to IPL. Run `xedit zpxe conf a` and create `ZPXE CONF` with the following content (replace the example data accordingly):

```
HOST cobbler.example.com
IPL 600
```

This connects the Cobbler server next time you log in to the z/VM guest. If an installation is scheduled on the Cobbler server, it will be executed. To schedule the installation, run the following command on the Cobbler server:

```
> sudo cobbler system edit --name ID ❶ --netboot-enabled 1 ❷ --profile PROFILENAME ❸
```

- ❶ z/VM user ID.
- ❷ Enable IPLing from the network.
- ❸ Name of an existing profile, see [Section 5.3.1.3.3, “Adjusting the profile”](#).

5.3.4.3 IPLing a KVM guest installation

To start the guest installation, you first need to start the VM Guest defined in [Section 5.3.3.3.1, “Create a virtual disk image”](#). Before you begin, ensure the kernel and initrd are available for IPL.

5.3.4.3.1 Preparing the installation source

Kernel and initrd of the installation system need to be copied to the VM Host Server to IPL the VM Guest into the installation system.

1. Log in to the KVM host and make sure you can connect to the remote host or device serving the installation source.
2. Copy the following two files from the installation source to `/var/lib/libvirt/images/`. If the data is served from a remote host, use `ftp`, `sftp`, or `scp` to transfer the files:

```
/boot/s390x/initrd  
/boot/s390x/cd.ikr
```

3. Rename the files on the KVM host:

```
> sudo cd /var/lib/libvirt/images/  
> sudo mv initrd s15-initrd.boot
```

```
> sudo mv cd.ikr s15-kernel.boot
```

5.3.4.3.2 IPL the VM Guest

To IPL the VM Guest, log in to the KVM host and run the following command:

```
> virsh create s15-1.xml --console
```

The installation process starts when the VM Guest is up and running, and you should see the following message:

```
Domain s15-1 started
Connected to domain s15-1
Escape character is ^]
Initializing cgroup subsys cpuset
Initializing cgroup subsys cpu
Initializing
cgroup subsys cpuacct
.
.
Please make sure your installation medium is available.
Retry?
0) <-- Back <--
1) Yes
2) No
```

Answer 2) *No* and choose *Installation* on the next step. Proceed as described in [Section 5.3.5.3, “Set up the network and select the installation source”](#).

5.3.5 Network configuration

Wait until the kernel has completed its start-up routines. If you perform the installation in basic mode or in an LPAR, open the *Operating System Messages* on the HMC or SE.

First, choose *Start Installation* in the `linuxrc` main menu. Then choose *Start Installation or Update* to start the installation process. Select *Network* as the installation medium, then select the type of network protocol to use for the installation. [Section 5.3.1, “Making the installation data available”](#) describes how to make the installation data available for the various types of network connections. Currently, *FTP*, *HTTP*, *NFS*, and *SMB/CIFS* (Windows file sharing) are supported.

From the list of available devices, choose an OSA or HiperSockets network device for receiving the installation data. Although the list may contain CTC, ESCON, or IUCV devices, they are no longer supported on SUSE Linux Enterprise Server.

5.3.5.1 Configure a HiperSockets interface

Select a HiperSocket device from the list of network devices. Then enter values for the read, write, and data channels:

EXAMPLE 5.5: SUPPORTED NETWORK CONNECTION TYPES AND DRIVER PARAMETERS

```
Choose the network device.

1) IBM parallel CTC Adapter (0.0.0600)
2) IBM parallel CTC Adapter (0.0.0601)
3) IBM parallel CTC Adapter (0.0.0602)
4) IBM Hipersocket (0.0.0800)
5) IBM Hipersocket (0.0.0801)
6) IBM Hipersocket (0.0.0802)
7) IBM OSA Express Network card (0.0.0700)
8) IBM OSA Express Network card (0.0.0701)
9) IBM OSA Express Network card (0.0.0702)
10) IBM OSA Express Network card (0.0.f400)
11) IBM OSA Express Network card (0.0.f401)
12) IBM OSA Express Network card (0.0.f402)
13) IBM IUCV

> 4

Device address for read channel. (Enter '+++' to abort).
[0.0.0800]> 0.0.0800

Device address for write channel. (Enter '+++' to abort).
[0.0.0801]> 0.0.0801

Device address for data channel. (Enter '+++' to abort).
[0.0.0802]> 0.0.0802
```

5.3.5.2 Configure an OSA express device

Select an OSA Express device from the list of network devices and specify a port number. Enter the values for the read, write and data channels. Choose whether to enable OSI Layer 2 support. The port number is required for the new 2 port OSA Express 3 Network devices. If you are not using an OSA Express 3 device, enter 0. OSA Express cards can also run in the “OSI layer 2 support” mode or the older more common “layer 3” mode. The card mode affects all systems that share the device, including systems on other LPARs. If in doubt, specify 2 for compatibility with the default mode used by other operating systems such as z/VM and z/OS. Consult with your hardware administrator for further information on these options.

EXAMPLE 5.6: NETWORK DEVICE DRIVER PARAMETERS

Choose the network device.

- 1) IBM parallel CTC Adapter (0.0.0600)
- 2) IBM parallel CTC Adapter (0.0.0601)
- 3) IBM parallel CTC Adapter (0.0.0602)
- 4) IBM Hipersocket (0.0.0800)
- 5) IBM Hipersocket (0.0.0801)
- 6) IBM Hipersocket (0.0.0802)
- 7) IBM OSA Express Network card (0.0.0700)
- 8) IBM OSA Express Network card (0.0.0701)
- 9) IBM OSA Express Network card (0.0.0702)
- 10) IBM OSA Express Network card (0.0.f400)
- 11) IBM OSA Express Network card (0.0.f401)
- 12) IBM OSA Express Network card (0.0.f402)
- 13) IBM IUCV

> 7

Enter the relative port number. (Enter '+++' to abort).

> 0

Device address for read channel. (Enter '+++' to abort).

[0.0.0700]> 0.0.0700

Device address for write channel. (Enter '+++' to abort).

[0.0.0701]> 0.0.0701

Device address for data channel. (Enter '+++' to abort).

[0.0.0702]> 0.0.0702

Enable OSI Layer 2 support?

0) <-- Back <--

1) Yes

2) No

> 1

MAC address. (Enter '+++' to abort).

> +++

5.3.5.3 Set up the network and select the installation source

After all network device parameters have been entered, the respective driver is installed and you see the corresponding kernel messages.

Next, you need to specify whether to use DHCP autoconfiguration for setting up the network interface parameters. Because DHCP only works on a few devices and requires special hardware configuration settings, choose *NO*. Doing this prompts you to specify the following networking parameters:

- The IP address of the system to install
- The corresponding netmask (if not having been specified with the IP address)
- The IP address of a gateway to reach the server
- A list of search domains covered by the domain name server (DNS)
- The IP address of your domain name server

EXAMPLE 5.7: NETWORKING PARAMETERS

```
Automatic configuration via DHCP?

0) <-- Back <--
1) Yes
2) No

> 2

Enter your IP address with network prefix.

You can enter more than one, separated by space, if necessary.
Leave empty for autoconfig.

Examples: 192.168.5.77/24 2001:db8:75:fff::3/64. (Enter '+++' to abort).
> 192.168.0.20/24

Enter your name server IP address.

You can enter more than one, separated by space, if necessary.
Leave empty if you don't need one.

Examples: 192.168.5.77 2001:db8:75:fff::3. (Enter '+++' to abort).
> 192.168.0.1

Enter your search domains, separated by a space:. (Enter '+++' to abort).
> example.com
```

```
Enter the IP address of your name server. Leave empty if you do not need one. (Enter '+++' to abort).
> 192.168.0.1
```

Finally, provide the required information about the installation server, such as the IP address, the directory containing the installation data, and login credentials. The installation system loads when the required information has been provided.

5.3.6 Connecting to the SUSE Linux Enterprise Server installation system

After loading the installation system, `linuxrc` prompts you to choose what type of display to use to control the installation procedure. The available options include Remote X11 (X Window System), VNC (Virtual Network Computing protocol), SSH (text mode or X11 installation via Secure Shell), Text-based UI and Graphical UI. The latter starts YaST in graphical mode on a local graphics display if it exists. On the s390x architecture, a local graphics display can be implemented using QEMU and the `virtio-gpu` driver.

The recommended options are VNC or SSH.

If the Text-based UI option is selected, YaST starts in text mode, and you can perform the installation directly within your terminal. See *Book "Administration Guide", Chapter 4 "YaST in text mode"* for instructions on how to use YaST in the text mode. The Text-based UI option is only useful when installing into LPAR.



Note: Terminal emulation for Text-based UI

To be able to work with YaST in the text mode, it needs to run in a terminal with VT220/Linux emulation (also called Text-based UI).

5.3.6.1 Initiating the installation for VNC

To remotely control an installation via VNC, follow these steps:

1. Selecting the VNC option starts the VNC server. A short note in the console displays the IP address and display number for connecting with `vncviewer`.
2. Enter the IP address and the display number of the SUSE Linux Enterprise Server installation system when prompted to do so.

3. When prompted, enter the IP address and the display number of the SUSE Linux Enterprise Server installation system.

```
http://<IP address of installation system>:5801/
```

4. After the connection has been established, install SUSE Linux Enterprise Server with YaST.

5.3.6.2 Initiating the installation for the X Window system



Important: X authentication mechanism

The direct installation with the X Window System relies on an authentication mechanism based on host names. This mechanism is disabled in current SUSE Linux Enterprise Server versions. We recommend performing the installation using SSH or VNC.

To remotely control an installation via X forwarding, follow these steps:

1. Make sure that the X server allows the client (the system that is installed) to connect. Set the variable **DISPLAYMANAGER_XSERVER_TCP_PORT_6000_OPEN="yes"** in the file `/etc/sysconfig/displaymanager`. Restart the X server and allow client binding to the server using **xhost CLIENT_IP_ADDRESS**.
2. When prompted at the installation system, enter the IP address of the machine running the X server.
3. Wait until YaST opens, then start the installation.

5.3.6.3 Initiating the installation for SSH

To connect to an installation system with the name `earth` via SSH, use the **ssh -X earth** command. If your workstation runs on Microsoft Windows, use the Putty tool available from <https://www.chiark.greenend.org.uk/~sgtatham/putty/>. Set *Enable X11 forwarding* in Putty under *Connection > SSH > X11*.

If you use another operating system, execute **ssh -X earth** to connect to an installation system with the name `earth`. X-Forwarding over SSH is supported if you have a local X server available. Otherwise, YaST provides a text interface over ncurses.

When prompted, enter the `root` user name and log in with your password. Enter `yast.ssh` to start YaST. YaST then guides you through the installation.

! Important: Fixing YaST over SSH issue

In certain situations, running the GUI version of YaST over SSH with X forwarding may fail with the following error message:

```
XIO: fatal IO error 11 (Resource temporarily unavailable) on X server
"localhost:11.0"
```

In this case you have two options.

- Run YaST with the `QT_XCB_GL_INTEGRATION=none` option, for example:

```
QT_XCB_GL_INTEGRATION=none yast.ssh
QT_XCB_GL_INTEGRATION=none yast2 disk
```

- Run the ncurses version of YaST application by disabling X forwarding or by specifying ncurses as the desired UI. To do the latter, use the `yast2 disk --ncurses` or `YUI_PREFERRED_BACKEND=ncurses yast2 disk` command.

Proceed with the installation procedure as described in [Chapter 9, Installation steps](#).

5.3.7 The SUSE Linux Enterprise Server boot procedure on IBM Z

On SLES 10 and 11 the boot process was handled by the `zipl` boot loader. To enable booting from Btrfs partitions and supporting system rollbacks with Snapper, the way SUSE Linux Enterprise Server is booted on IBM Z has changed.

GRUB 2 replaces `zipl` on SUSE Linux Enterprise Server for IBM Z. GRUB 2 on the AMD64/Intel 64 architecture includes device drivers on the firmware level to access the file system. On the mainframe there is no firmware and adding `ccw` to GRUB 2 would not only be a major undertaking, but would also require a reimplementation of `zipl` in GRUB 2. Therefore SUSE Linux Enterprise Server uses a two-stage approach:

Stage one:

A separate partition containing the kernel and an `initrd` is mounted to `/boot/zipl`. This kernel and the `initrd` are loaded via `zipl` using the configuration from `/boot/zipl/config`.

This configuration adds the keyword `initgrub` to the kernel command line. When the kernel and `initrd` are loaded, the `initrd` activates the devices required to mount the root file system (see `/boot/zipl/active_devices.txt`). Afterward a GRUB 2 user space program is started, which reads `/boot/grub2/grub.cfg`.

Stage two:

The kernel and the `initrd` specified in `/boot/grub2/grub.cfg` are started via `kexec`. Devices listed in `/boot/zipl/active_devices.txt` that are necessary for starting the on-disk system are then activated. Other devices from that list will be whitelisted, but otherwise ignored. The root file system is mounted and the boot procedure continues like on the other architectures.

For more details on the boot process, refer to *Book “Administration Guide”, Chapter 16 “Introduction to the boot process”*.

5.4 Secure boot

For the secure boot functionality to work on a IBM Z system, the following conditions must be met.

- The machine must be z15 T01, z15 T02, LinuxONE III LT1, LinuxONE III LT2, or a later model.
- You must use an LPAR (secure boot is not supported on z/VM and KVM).
- The LPAR must have secure boot enabled.
- You must use SCSI (FCP) disks (secure boot is not supported on DASD).



Note: Hardware migration

In case you migrate to a different machine (for example, from z13 to z15), ensure that the LPAR on the target machine has the secure boot state of the system on its disk.

Changing the secure boot state must be performed according to the following procedure.

PROCEDURE 5.4: CHANGING SECURE BOOT STATE

1. Enable secure boot in YaST and write the new boot loader.

2. Shut down the system.
3. Change the configuration of the LPAR (enable or disable secure boot).
4. Boot the system.



Note: Secure boot on HMC

The system on the disk configured with the `secure=1` parameter can be booted on z15 HMC as long as the firmware supports the new on-disk format (which is always the case on z15).

5.5 The parmfile—automating the system configuration

The installation process can be partially automated by specifying the essential parameters in the `parmfile`. The `parmfile` contains all the data required for network setup and DASD configuration. In addition to that, it can be used to set up the connection method to the SUSE Linux Enterprise Server installation system and the YaST instance running there. This reduces user interaction to the actual YaST installation.

The parameters listed in [Section 5.5.1, “General parameters”](#) can be passed to the installation routine as the default values for installation. Note that all IP addresses, server names, and numerical values are examples. Replace them with the actual values of your installation scenario.

The number of lines in the `parmfile` is limited to 10. You can specify more than one parameter on a line. Parameter names are not case-sensitive. Parameters must be separated by spaces. You may specify the parameters in any order. Always keep the `PARAMETER=value` string together on one line. The length of each line must not exceed 80 characters. For example:

```
Hostname=s390zvm01.suse.de HostIP=10.11.134.65
```



Tip: Using IPv6 during the installation

By default, you can only assign IPv4 network addresses to your machine. To enable IPv6 during installation, specify one of the following parameters at the boot prompt: `ipv6=1` (accept IPv4 and IPv6) or `ipv6only=1` (accept IPv6 only).

Some of the following parameters are required. If they are missing, the automatic process prompts you to specify them.

5.5.1 General parameters

AutoYaST=<URL> Manual=0

The AutoYaST parameter specifies the location of the autoinst.xml control file for automatic installation. The Manual parameter controls if the other parameters are only default values that still must be acknowledged by the user. Set this parameter to 0 if all values should be accepted and no questions asked. Setting AutoYaST defaults Manual to 0.

DeviceAutoConfig=<0|1|2>

In linuxrc, the DeviceAutoConfig parameter controls the use of I/O device auto-configuration data for IBM Z systems.

If set to 0, auto-configuration is disabled. If set to 1, the existing auto-config data are applied. If set to 2 (the default), a dialog is shown if auto-config data are present. The user is asked whether to apply them.

For more details, see [Section 5.5.4, “I/O device auto-configuration on IBM Z systems”](#).

Info=<URL>

Specifies a location for a file with additional options. This helps to overcome the limitations of 10 lines (and 80 characters per line under z/VM) for the parmfile. Further documentation on the Info file can be found in *Book “AutoYaST Guide”, Chapter 9 “The auto-installation process”, Section 9.3.3 “Combining the **linuxrc**info file with the AutoYaST control file”*. Since the Info file can typically only be accessed through the network on IBM Z, you cannot use it to specify the options required to set up the network (that is, the options described in [Section 5.5.2, “Configuring the network interface”](#)). Other linuxrc-specific options, such as those related to debugging, must be specified in the parmfile itself.

Upgrade=<0|1>

To upgrade SUSE Linux Enterprise, specify Upgrade=1. A custom parmfile is required for upgrading an existing installation of SUSE Linux Enterprise. Without this parameter, the installation provides no upgrade option.

5.5.2 Configuring the network interface



Important: Configuring the network interface

The settings described in this section apply only to the network interface used during installation. Configure additional network interfaces in the installed system by following the instructions in *Book "Administration Guide", Chapter 23 "Basic networking", Section 23.5 "Configuring a network connection manually"*.

Hostname=zsystems.example.com

Enter the fully qualified host name.

Domain=example.com

Domain search path for DNS. Allows you to use short host names instead of fully qualified ones.

HostIP=192.168.1.2/24

Enter the IP address of the interface to configure.

Gateway=192.168.1.3

Specify the gateway to use.

Nameserver=192.168.1.4

Specify the DNS server in charge.

InstNetDev=osa

Enter the type of interface to configure. Possible values are osa, hsi, ctc, escon, and iucv (CTC, ESCON, and IUCV are no longer supported).

For the ctc interfaces escon and iucv (CTC, ESCON, and IUCV are no longer supported), enter the IP address of the peer:

```
Pointopoint=192.168.55.20
```

OsaInterface=<lcs|qdio>

For osa network devices, specify the host interface (qdio or lcs).

Layer2=<0|1>

For osa QDIO Ethernet and hsi devices, specify whether to enable (1) or disable (0) OSI Layer 2 support.

OSAHWAddr=02:00:65:00:01:09

For Layer 2-enabled osa QDIO Ethernet devices. Either specify a MAC address manually or state OSAHWADDR= (with trailing white space) for the system default.

PortNo=<0|1>

For osa network devices, specify the port number (provided the device supports this feature). The default value is 0.

Each of the interfaces requires certain setup options:

- Interfaces ctc and escon (CTC and ESCON are no longer supported):

```
ReadChannel=0.0.0600
WriteChannel=0.0.0601
```

ReadChannel specifies the READ channel to use. WriteChannel specifies the WRITE channel.

- For the ctc interface (no longer supported), specify the protocol that should be used for this interface:

```
CTCProtocol=<0/1/2>
```

Valid entries would be:

<u>0</u>	Compatibility mode, also for non-Linux peers other than OS/390 and z/OS (this is the default mode)
<u>1</u>	Extended mode
<u>2</u>	Compatibility mode with OS/390 and z/OS

- Network device type osa with interface lcs:

```
ReadChannel=0.0.0124
```

ReadChannel stands for the channel number used in this setup. A second port number can be derived from this by adding one to ReadChannel. Portnumber is used to specify the relative port.

- Interface iucv:

```
IUCVPeer=PEER
```

Enter the name of the peer machine.

- Network device type osa with interface qdio for OSA-Express Gigabit Ethernet:

```
ReadChannel=0.0.0700
WriteChannel=0.0.0701
DataChannel=0.0.0702
```

For ReadChannel, enter the number of the READ channel. For WriteChannel, enter the number of the WRITE channel. DataChannel specifies the DATA channel. Make sure that the READ channel has an even device number.

- Interface hsi for HiperSockets and VM guest LANs:

```
ReadChannel=0.0.0800
WriteChannel=0.0.0801
DataChannel=0.0.0802
```

For ReadChannel, enter the appropriate number for the READ channel. For WriteChannel and DataChannel, enter the WRITE and DATA channel numbers.

5.5.3 Specifying the installation source and YaST interface

Install=nfs://server/directory/DVD1/

Specify the location of the installation source to use. Supported protocols are nfs, smb (Samba/CIFS), ftp, tftp, http, and https.

If an ftp, tftp or smb URL is provided, specify the user name and password. Skip credentials for anonymous or guest login.

```
Install=ftp://USER:PASSWORD@SERVER/DIRECTORY/DVD1/
Install=tftp://USER:PASSWORD@SERVER/DIRECTORY/DVD1/
```

If you want to perform the installation over an encrypted connection, use an https URL. If the certificate cannot be verified, use the sslcerts=0 boot option to disable certificate checking.

In case of a Samba or CIFS installation, you can also specify the domain:

```
Install=smb://WORKDOMAIN;USER:PASSWORD@SERVER/DIRECTORY/DVD1/
```


ssh=1 vnc=1 Display_IP=192.168.42.42

The installation method depends on which parameter you specify. ssh enables SSH installation, vnc starts a VNC server on the installing machine, and Display_IP causes the installing system to try to connect to an X server at the specified address. Only one of these parameters should be set.



Important: X authentication mechanism

The direct installation with the X Window System relies on an authentication mechanism based on host names. This mechanism is disabled on current SUSE Linux Enterprise Server versions. We recommend to perform an installation using SSH or VNC is preferred.

To allow a connection between YaST and the remote X server, run **xhost <IP address>** with the address of the installing machine on the remote machine.

For VNC, specify a password of six to eight characters to use for installation:

```
VNCPassword=<a password>
```

For SSH, specify a password of six to eight characters to use for installation:

```
ssh.password=<a password>
```

5.5.4 I/O device auto-configuration on IBM Z systems

I/O device auto-configuration is a mechanism that allows users to specify IDs and settings of I/O devices that should be automatically enabled in Linux. This information is specified for an LPAR via an HMC running in DPM (Dynamic Partition Manager) mode.



Note

The I/O device auto-configuration functionality is available on systems with the DPM running. DPM runs by default on LinuxONE machines. For IBM Z, this functionality must be ordered.

In **linuxrc**, the DeviceAutoConfig parameter controls the use of I/O device auto-configuration data for IBM Z systems.

DeviceAutoConfig=0

If set to 0, auto-configuration is disabled.

DeviceAutoConfig=1

If set to 1, existing auto-config data are applied.

DeviceAutoConfig=2 (default)

If set to 2 (the default), a dialog is shown if auto-config data are present. The user is asked whether to apply them.

If device auto-config is disabled by the user, the kernel parameter `rd.zdev=no-auto` is added to the boot options of the target system.

To enable I/O auto-configuration using YaST, run the **yast2 system_settings** command, switch to the *Kernel Settings* section, and enable the *Enable I/O device auto-configuration* option. To disable I/O auto-configuration in an AutoYaST profile, add the following kernel parameter in the append section of the global boot loader options. For example:

```
<bootloader>
  <global>
    <append>rd.zdev=no-auto</append>
  </global>
</bootloader>
```

For more context on the AutoYaST boot loader options, see *Book "AutoYaST Guide", Chapter 4 "Configuration and installation options", Section 4.4 "The GRUB 2 boot loader"*.

During installation, the status of the auto-configuration setting is displayed in the *Device Settings* section of the *Installation Settings* screen.

5.5.5 Example parmfiles

The maximum capacity of a parmfile is 860 characters. As a rule of thumb, the parmfile should contain a maximum of 10 lines with no more than 79 characters. When reading a parmfile, all lines are concatenated without adding white spaces, therefore the last character (79) of each line needs to be a **Space** .

To receive potential error messages on the console, use

```
linuxrclog=/dev/console
```

EXAMPLE 5.8: PARMFILE FOR AN INSTALLATION FROM NFS WITH VNC AND AUTOYAST, WITH I/O DEVICE AUTO CONFIGURATION

```
ramdisk_size=131072 root=/dev/ram1 ro init=/linuxrc TERM=dumb
instnetdev=osa osainterface=qdio layer2=1 osahwaddr=
pointopoint=192.168.0.1 hostip=192.168.0.2
nameserver=192.168.0.3 DeviceAutoConfig=1
install=nfs://192.168.0.4/SLES/SLES-12-Server/s390x/DVD1
autoyast=http://192.168.0.5/autoinst.xml
linuxrclog=/dev/console vnc=1 VNCPassword=testing
```

EXAMPLE 5.9: PARMFILE FOR INSTALLATION WITH NFS, SSH, AND HSI AND AUTOYAST WITH NFS

```
ramdisk_size=131072 root=/dev/ram1 ro init=/linuxrc TERM=dumb
AutoYast=nfs://192.168.1.1/autoinst/s390.xml
Hostname=zsystems.example.com HostIP=192.168.1.2
Gateway=192.168.1.3 Nameserver=192.168.1.4
InstNetDev=hsi layer2=0
Netmask=255.255.255.128 Broadcast=192.168.1.255
readchannel=0.0.702c writechannel=0.0.702d datachannel=0.0.702e
install=nfs://192.168.1.5/SLES-12-Server/s390x/DVD1/
ssh=1 ssh.password=testing linuxrclog=/dev/console
```

EXAMPLE 5.10: PARMFILE FOR INSTALLATION IN VLAN

```
ro ramdisk_size=50000 MANUAL=0 PORTNO=1 ReadChannel=0.0.b140
WriteChannel=0.0.b141 DataChannel=0.0.b142
cio_ignore=all,!condev,!0.0.b140-0.0.b142,!0.0.e92c,!0.0.5000,!0.0.5040
HostIP= Gateway= Hostname=zsystems.example.com nameserver=192.168.0.1
Install=ftp://user:password@10.0.0.1/s390x/SLES15.0/INST/ usevnc=1
vncpassword=12345 InstNetDev=osa Layer2=1 OSAInterface=qdio ssl_certs=0
osahwaddr= domain=example.com self_update=0
vlanid=201
```

5.6 Using the vt220 terminal emulator

Recent MicroCode Levels allow the use of an integrated vt220 terminal emulator (ASCII terminal) in addition to the standard line mode terminal. The vt220 terminal is connected to `/dev/ttysclp0`. The line mode terminal is connected to `/dev/ttysclp_line0`. For LPAR installations, the vt220 terminal emulator is activated by default.

To start the Text-based UI on HMC, log in to the HMC, and select *Systems Management > Systems > IMAGE_ID* . Select the radio button for the LPAR and select *Recovery > Integrated ASCII Console*.

To redirect the kernel messages at boot time from the system console to the vt220 terminal, add the following entries to the `parameters` line in `/etc/zipl.conf`:

```
console=ttysclp0 console=ttysclp_line0
```

The resulting `parameters` line would look like the following example:

```
parameters = "root=/dev/dasda2 TERM=dumb console=ttysclp0 console=ttysclp_line0"
```

Save the changes in `/etc/zipl.conf`, run **zipl**, and reboot the system.

5.7 More information

Find further technical documentation about IBM Z in the IBM Redbooks (<https://www.redbooks.ibm.com/Redbooks.nsf/domains/zsystems>) or at IBM developerWorks (<https://developer.ibm.com/>). SUSE Linux Enterprise Server-specific documentation is available from <https://developer.ibm.com/technologies/linux/>.

5.7.1 General documents about Linux on IBM Z

A general coverage of Linux on IBM Z can be found in the following documents:

- Linux on IBM eServer zSeries and S/390: ISP and ASP Solutions (SG24-6299)


These documents might not reflect the current state of Linux, but the principles of Linux deployment outlined there remain accurate.

5.7.2 Technical issues of Linux on IBM Z

Refer to the following documents for technical information about the Linux kernel and application topics. For the most recent versions of the documents, visit (<https://developer.ibm.com/technologies/linux/>).

- Linux on System z Device Drivers, Features, and Commands
- zSeries ELF Application Binary Interface Supplement
- Linux on System z Device Drivers, Using the Dump Tools



- IBM zEnterprise 196 Technical Guide
- IBM zEnterprise EC12 Technical Guide
- IBM z13 Technical Guide
- IBM z14 Technical Guide
- IBM z15 Technical Guide

A Redbook for Linux application development is available at <https://www.redbooks.ibm.com> :

- Linux on IBM eServer zSeries and S/390: Application Development (SG24-6807)

5.7.3 Advanced configurations for Linux on IBM Z

Refer to the following Redbooks, Redpapers, and online resources for more complex IBM Z scenarios:

- Linux on IBM eServer zSeries and S/390: Large Scale Deployment (SG24-6824)
- Linux on IBM eServer zSeries and S/390: Performance Measuring and Tuning (SG24-6926)
- Linux with zSeries and ESS: Essentials (SG24-7025)
- IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services with IBM eServer zSeries (SG24-5680)
- Linux on IBM zSeries and S/390: High Availability for z/VM and Linux (REDP-0220)
- Saved Segments Planning and Administration
<https://publibfp.boulder.ibm.com/epubs/pdf/hcsg4c10.pdf> 
- Linux on System z documentation for "Development stream"
- Introducing IBM Secure Execution for Linux, Securing the guest
https://www.ibm.com/support/knowledgecenter/linuxonibm/com.ibm.linux.z.lxse/lxse_t_secureexecution.html 

6 Installation on virtualization hosts

This section describes the support status of SUSE Linux Enterprise Server 15 SP6 running as a guest operating system on top of different virtualization hosts (hypervisors).

TABLE 6.1: THE FOLLOWING SUSE HOST ENVIRONMENTS ARE SUPPORTED

SUSE Linux Enterprise Server	Hypervisors
SUSE Linux Enterprise Server 11 SP4	Xen and KVM
SUSE Linux Enterprise Server 12 SP1 to SP5	Xen and KVM
SUSE Linux Enterprise Server 15 GA to SP6	Xen and KVM

THE FOLLOWING THIRD-PARTY HOST ENVIRONMENTS ARE SUPPORTED

- Citrix XenServer (<https://www.citrix.com/products/citrix-hypervisor/>) ↗
- Hyper-V (<https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/supported-suse-virtual-machines-on-hyper-v>) ↗
- Nutanix Acropolis Hypervisor with AOS (<https://www.nutanix.com/products>) ↗
- Oracle VM Server 3.4 (<https://www.oracle.com/fr/virtualization/virtualbox/>) ↗
- Oracle Linux KVM 7, 8 (<https://www.oracle.com/linux/>) ↗
- VMware ESXi 6.5, 6.7, 7.0, 8.0 (<https://www.vmware.com/products/esxi-and-esx.html>) ↗
- Windows Server 2022

You can also search in the SUSE YES certification database (<https://www.suse.com/yessearch/Search.jsp>) ↗

THE LEVEL OF SUPPORT IS AS FOLLOWS

- Support for SUSE host operating systems is full L3 (both for the guest and host) according to the respective [product life cycle](https://www.suse.com/lifecycle/) (<https://www.suse.com/lifecycle/>) ↗.
- SUSE provides full L3 support for SUSE Linux Enterprise Server guests within third-party host environments.
- Support for the host and cooperation with SUSE Linux Enterprise Server guests must be provided by the host system's vendor.

7 Installation on hardware not supported at release

With some newer hardware, the installation medium of SUSE Linux Enterprise Server cannot boot. This can be the case when the hardware did not exist at the time of the release of SUSE Linux Enterprise Server. For such a situation SUSE provides *Kernel Update ISO (kISO)* images. This chapter describes how to use the kernel update to install SUSE Linux Enterprise Server on current hardware.

7.1 Download kernel update

Kernel Update ISO images are available on the SUSE SolidDriver home page. Use <https://drivers.suse.com> to search for bootable ISO images for your vendor and operating system version.

You can download the full ISO image or only the `initrd` and `linux` files. The ISO usually needs to be copied to a USB flash drive or burned to a DVD. The `initrd` and `linux` files can be used for a PXE boot. For details about booting via PXE, see [Chapter 18, Preparing network boot environment](#).

7.2 Boot kernel update

To use the kernel update, boot from the USB flash drive or via PXE. When the `linux` and the `initrd` are loaded, you will be asked to insert the installation medium.

You can use the boot parameters described in [Chapter 8, Boot parameters](#). This allows using other installation sources than the installation USB flash drive.

II Installation procedure

- 8 Boot parameters **79**
- 9 Installation steps **97**
- 10 Registering SUSE Linux Enterprise and managing modules/extensions **146**
- 11 *Expert Partitioner* **157**
- 12 Remote installation **181**
- 13 Troubleshooting **189**

8 Boot parameters

SUSE Linux Enterprise Server allows setting several parameters during boot, for example choosing the source of the installation data or setting the network configuration.

Using the appropriate set of boot parameters helps simplify your installation procedure. Many parameters can also be configured later using the `linuxrc` routines, but using the boot parameters is easier. In some automated setups, the boot parameters can be provided with `initrd` or an `info` file.

The way the system is started for the installation depends on the architecture—system start-up is different for PC (AMD64/Intel 64) or mainframe, for example. If you install SUSE Linux Enterprise Server as a VM Guest on a KVM or Xen hypervisor, follow the instructions for the AMD64/Intel 64 architecture.



Note: Boot options and boot parameters

The terms *Boot Parameters* and *Boot Options* are often used interchangeably. In this documentation, we mostly use the term *Boot Parameters*.

8.1 Using the default boot parameters

The boot parameters are described in detail in [Chapter 9, Installation steps](#). Generally, selecting *Installation* starts the installation boot process.

If problems occur, use *Installation—ACPI Disabled* or *Installation—Safe Settings*. For more information about troubleshooting the installation process, refer to [Chapter 13, Troubleshooting](#).

The menu bar at the bottom of the screen offers some advanced functionality needed in some setups. Using the function keys (`F1` ... `F12`), you can specify additional options to pass to the installation routines without having to know the detailed syntax of these parameters (see [Chapter 8, Boot parameters](#)). A detailed description of the available function keys is available in [Section 8.2.1, “The boot screen on machines with traditional BIOS”](#).

8.2 PC (AMD64/Intel 64/AArch64)

This section describes changing the boot parameters for AMD64, Intel 64 and AArch64.

8.2.1 The boot screen on machines with traditional BIOS

The boot screen displays several options for the installation procedure. *Boot from Hard Disk* boots the installed system and is selected by default. Select one of the other options with the arrow keys and press **Enter** to boot it. The relevant options are:

Installation

The normal installation mode. All modern hardware functions are enabled. In case the installation fails, see **F5** *Kernel* for boot parameters that disable potentially problematic functions.

Upgrade

Perform a system upgrade. For more information refer to Book “Upgrade Guide”, Chapter 2 “Upgrade paths and methods”.

More > Rescue System

Starts a minimal Linux system without a graphical user interface.

More > Boot Linux System

Boot a Linux system that is already installed. You will be asked from which partition to boot the system.

More > Check Installation Media

This option is only available when you install from media created from downloaded ISOs. In this case it is recommended to check the integrity of the installation medium. This option starts the installation system before automatically checking the media. In case the check was successful, the normal installation routine starts. If a corrupt media is detected, the installation routine aborts. Replace the broken medium and restart the installation process.

More > Memory Test

Tests your system RAM using repeated read and write cycles. Terminate the test by rebooting. For more information, see [Section 13.4, “Boot failure”](#).

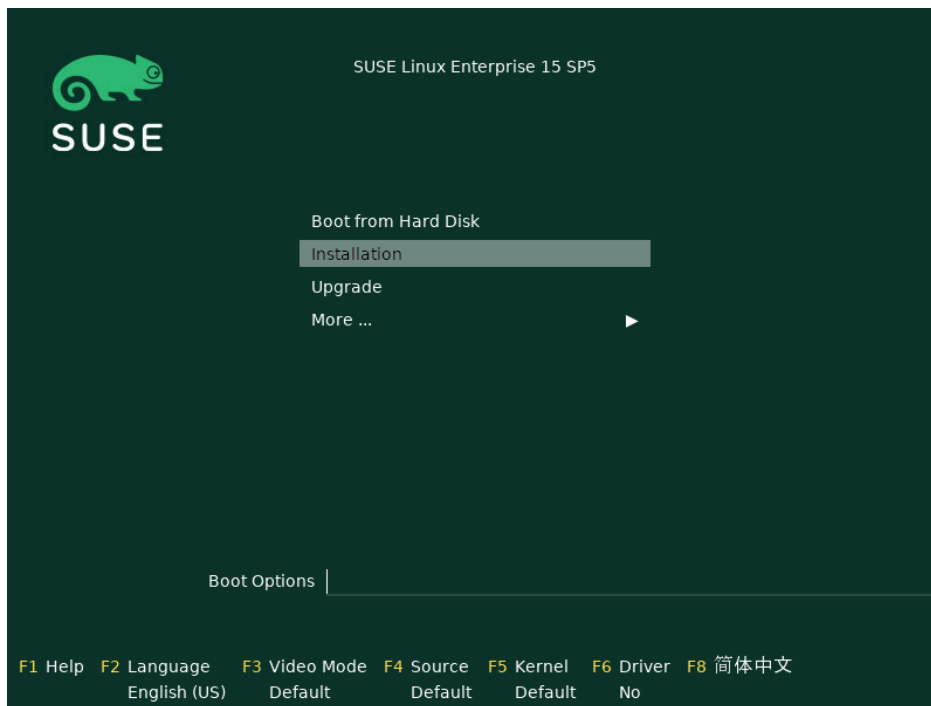


FIGURE 8.1: THE BOOT SCREEN ON MACHINES WITH A TRADITIONAL BIOS

Use the function keys shown at the bottom of the screen to change the language, screen resolution, installation source or to add an additional driver from your hardware vendor:

F1 *Help*

Get context-sensitive help for the active element of the boot screen. Use the arrow keys to navigate, **Enter** to follow a link, and **Esc** to leave the help screen.

F2 *Language*

Select the display language and a corresponding keyboard layout for the installation. The default language is English (US).

F3 *Video Mode*

Select various graphical display modes for the installation. By *Default* the video resolution is automatically determined using KMS (“Kernel Mode Setting”). If this setting does not work on your system, choose *No KMS* and, optionally, specify vga=ask on the boot command line to get prompted for the video resolution. Choose *Text Mode* if the graphical installation causes problems.

F4 Source

Normally, the installation is performed from the inserted installation medium. Here, select other sources, like FTP or NFS servers, or configure a proxy server. If the installation is deployed on a network with an SLP server, select an installation source available on the server with this option. Find information about setting up an installation server with SLP at [Chapter 17, Setting up a network installation source](#).

F5 Kernel

If you encounter problems with the regular installation, this menu offers to disable a few potentially problematic functions. If your hardware does not support ACPI (advanced configuration and power interface) select *No ACPI* to install without ACPI support. *No local APIC* disables support for APIC (Advanced Programmable Interrupt Controllers) which may cause problems with some hardware. *Safe Settings* boots the system with the DMA mode (for CD/DVD-ROM drives) and power management functions disabled.

If you are not sure, try the following options first: *Installation—ACPI Disabled* or *Installation—Safe Settings*. Experts can also use the command line (*Boot Options*) to enter or change kernel parameters.

F6 Driver

Press this key to notify the system that you have an optional driver update for SUSE Linux Enterprise Server. With *File* or *URL*, load drivers directly before the installation starts. If you select *Yes*, you are prompted to insert the update disk at the appropriate point in the installation process.



Tip: Getting driver update disks

Driver updates for SUSE Linux Enterprise are provided at <https://driver-s.suse.com/>. These drivers have been created via the SUSE SolidDriver Program.

8.2.2 The boot screen on machines equipped with UEFI

UEFI (Unified Extensible Firmware Interface) is a new industry standard which replaces and extends the traditional BIOS. The latest UEFI implementations contain the “Secure Boot” extension, which prevents booting malicious code by only allowing signed boot loaders to be executed. See Book “Administration Guide”, Chapter 17 “UEFI (Unified Extensible Firmware Interface)” for more information.

The boot manager GRUB 2, used to boot machines with a traditional BIOS, does not support UEFI, therefore GRUB 2 is replaced with GRUB 2 for EFI. If Secure Boot is enabled, YaST will automatically select GRUB 2 for EFI for installation. From an administrative and user perspective, both boot manager implementations behave the same and are called GRUB 2 in the following.



Tip: Using additional drivers with Secure Boot

When installing with Secure Boot enabled, you cannot load drivers that are not shipped with SUSE Linux Enterprise Server. This is also true of drivers shipped via SolidDriver, because their signing key is not trusted by default.

To load drivers not shipped with SUSE Linux Enterprise Server, do either of the following:

- Before the installation, add the needed keys to the firmware database via firmware/system management tools.
- Use a bootable ISO that will enroll the needed keys in the MOK list on the first boot.

For more information, see *Book "Administration Guide", Chapter 17 "UEFI (Unified Extensible Firmware Interface)", Section 17.1 "Secure boot"*.

The boot screen displays several options for the installation procedure. Change the selected option with the arrow keys and press **Enter** to boot it. The relevant options are:

Installation

The normal installation mode. All modern hardware functions are enabled. In case the installation fails, see **F5** *Kernel* for boot parameters that disable potentially problematic functions.

Upgrade

Perform a system upgrade. For more information refer to *Book "Upgrade Guide", Chapter 2 "Upgrade paths and methods"*.

More > Rescue System

Starts a minimal Linux system without a graphical user interface.

More > Boot Linux System

Boot a Linux system that is already installed. You will be asked from which partition to boot the system.

More > Check Installation Media

This option is only available when you install from media created from downloaded ISOs. In this case it is recommended to check the integrity of the installation medium. This option starts the installation system before automatically checking the media. In case the check was successful, the normal installation routine starts. If a corrupt media is detected, the installation routine aborts.




FIGURE 8.2: THE BOOT SCREEN ON MACHINES WITH UEFI

GRUB 2 for EFI on SUSE Linux Enterprise Server does not support a boot prompt or function keys for adding boot parameters. By default, the installation will be started with American English and the boot media as the installation source. A DHCP lookup will be performed to configure the network. To change these defaults or to add boot parameters you need to edit the respective boot entry. Highlight it using the arrow keys and press **E**. See the on-screen help for editing hints (note that only an English keyboard is available now). The *Installation* entry will look similar to the following:

```
setparams 'Installation'

set gfxpayload=keep
echo 'Loading kernel ...'
linuxefi /boot/x86_64/loader/linux splash=silent
echo 'Loading initial ramdisk ...'
initrdefi /boot/x86_64/loader/initrd
```

Add space-separated parameters to the end of the line starting with `linuxefi`. To boot the edited entry, press **F10**. If you access the machine via serial console, press **Esc** **0**. A complete list of parameters is available at <https://en.opensuse.org/Linuxrc>.

8.3 List of important boot parameters

This section contains a selection of important boot parameters.

8.3.1 General boot parameters

autoyast=URL

The autoyast parameter specifies the location of the autoinst.xml control file for automatic installation.

manual=<0 | 1>

The manual parameter controls whether the other parameters are only default values that still must be acknowledged by the user. Set this parameter to 0 if all values should be accepted and no questions asked. Setting autoyast implies setting manual to 0.

Info=URL

Specifies a location for a file from which to read additional options.

IBM Z This helps to overcome the limitations of 10 lines (and 80 characters per line under z/VM) for the parmfile. More documentation on the Info file can be found in *Book "AutoYaST Guide", Chapter 9 "The auto-installation process", Section 9.3.3 "Combining the **linuxrc** info file with the AutoYaST control file"*. Since the Info file can typically only be accessed through the network on IBM Z, you cannot use it to specify options required to set up the network (these options are described in [Section 8.3.2, "Configuring the network interface"](#)). Also other linuxrc specific options such as for debugging need to be specified in the parmfile to be effective. ◁

upgrade=<0 | 1>

To upgrade SUSE Linux Enterprise Server, specify Upgrade=1.

IBM Z A custom parmfile is required for upgrading an existing installation of SUSE Linux Enterprise. Without this parameter, the installation provides no upgrade option. ◁

dud=URL

Load driver updates from URL.

Set `dud=ftp://ftp.example.com/PATH_TO_DRIVER` or `dud=http://www.example.com/PATH_TO_DRIVER` to load drivers from a URL. When `dud=1` you will be asked for the URL during boot.

`language=LANGUAGE`

Set the installation language. Some supported values are `cs_CZ`, `de_DE`, `es_ES`, `fr_FR`, `ja_JP`, `pt_BR`, `pt_PT`, `ru_RU`, `zh_CN`, and `zh_TW`.

`acpi=off`

Disable ACPI support.

`noapic`

No logical APIC.

`nomodeset`

Disable KMS.

`textmode=1`

Start installer in text mode.

`console=SERIAL_DEVICE[,MODE]`

`SERIAL_DEVICE` can be an actual serial or parallel device (for example `ttyS0`) or a virtual terminal (for example `tty1`). `MODE` is the baud rate, parity and stop bit (for example `9600n8`). The default for this setting is set by the mainboard firmware. If you do not see output on your monitor, try setting `console=tty1`. It is possible to define multiple devices.

8.3.2 Configuring the network interface



Important: Configuring the network interface

The settings discussed in this section apply only to the network interface used during installation. Configure additional network interfaces in the installed system by following the instructions in *Book "Administration Guide", Chapter 23 "Basic networking", Section 23.5 "Configuring a network connection manually"*.

The network will only be configured if it is required during the installation. To force the network to be configured, use the `netsetup` or `ifcfg` parameters.

netsetup=VALUE

netsetup=dhcp forces a configuration via DHCP. Set netsetup=-dhcp when configuring the network with the boot parameters hostip, gateway and nameserver. With the option netsetup=hostip,netmask,gateway,nameserver the installer asks for the network settings during boot.

ifcfg=INTERFACE[.VLAN]=[.try,]SETTINGS

INTERFACE can be * to match all interfaces or, for example, eth* to match all interfaces that start with eth. It is also possible to use MAC addresses as values.

Optionally, a VLAN can be set behind the interface name, separated by a period.

If SETTINGS is dhcp, all matching interfaces will be configured with DHCP. If you add the try option, configuration will stop when the installation repository can be reached via one of the configured interfaces.

Alternatively, you can use static configuration. With static parameters, only the first matching interface will be configured, unless you add the try option. This will configure all interfaces until the repository can be reached.

The syntax for the static configuration is:

```
ifcfg=*="IPS_NETMASK,GATEWAYS,NAMESERVERS,DOMAINS"
```

Each comma separated value can in turn contain a list of space character separated values. IPS_NETMASK is in the *CIDR notation*, for example 10.0.0.1/24. The quotes are only needed when using space character separated lists. Example with two name servers:

```
ifcfg=*="10.0.0.10/24,10.0.0.1,10.0.0.1 10.0.0.2,example.com"
```



Tip: Other networking parameters

The ifcfg boot parameter is very powerful and allows you to set almost all networking parameters. In addition to the parameters mentioned above, you can set values for all configuration options (comma separated) from /etc/sysconfig/network/ifcfg.template and /etc/sysconfig/network/config. The following example sets a custom MTU size on an interface otherwise configured via DHCP:

```
ifcfg=eth0=dhcp,MTU=1500
```

hostname=host.example.com

Enter the fully qualified host name.

domain=example.com

Domain search path for DNS. Allows you to use short host names instead of fully qualified ones.

hostip=192.168.1.2[/24]

Enter the IP address of the interface to configure. The IP can contain the subnet mask, for example hostip=192.168.1.2/24. This setting is only evaluated if the network is required during the installation.

gateway=192.168.1.3

Specify the gateway to use. This setting is only evaluated if the network is required during the installation.

nameserver=192.168.1.4

Specify the DNS server in charge. This setting is only evaluated if the network is required during the installation.

domain=example.com

Domain search path. This setting is only evaluated if the network is required during the installation.

8.3.3 Specifying the installation source

If you are not using DVD or USB flash drive for installation, specify an alternative installation source.

install=SOURCE

Specify the location of the installation source to use. Possible protocols are cd, hd, slp, nfs, smb (Samba/CIFS), ftp, tftp, http, and https. Not all source types are available on all platforms. For example IBM Z does not support cd and hd. The default option is cd.

To install over an encrypted connection, use an https URL. If the certificate cannot be verified, disable certificate checking with the sslcerts=0 boot parameter.

If an http, https, ftp, tftp, or smb URL is given, you can authenticate by specifying the user name and password with the URL. Example:

```
install=https://USER:PASSWORD@SERVER/DIRECTORY/DVD1/
```

In case of a Samba or CIFS installation, you can also specify the domain that should be used:

```
install=smb://WORKDOMAIN;USER:PASSWORD@SERVER/DIRECTORY/DVD1/
```

To use `cd`, `hd` or `slp`, set them as the following example:

```
install=cd:/
install=hd:/?device=sda/PATH_TO_ISO
install=slp:/
```

8.3.4 Specifying remote access

Only one of the different remote control methods should be specified at a time. The different methods are: SSH, VNC, remote X server. For information about how to use the parameters listed in this section, see [Chapter 12, Remote installation](#).

`display_ip=IP_ADDRESS`

`Display_IP` causes the installing system to try to connect to an X server at the given address.



Important: X authentication mechanism

The direct installation with the X Window System relies on a primitive authentication mechanism based on host names. This mechanism is disabled on current SUSE Linux Enterprise Server versions. Installation with SSH or VNC is preferred.

`vnc=1`

Enables a VNC server during the installation.

`vncpassword=PASSWORD`

Sets the password for the VNC server.

`ssh=1`

`ssh` enables SSH installation.

`ssh.password=PASSWORD`

Specifies an SSH password for the root user during installation.

8.4 Advanced setups

To configure access to a local RMT or **supportconfig** server for the installation, you can specify boot parameters to set up these services during installation. The same applies if you need IPv6 support during the installation.

8.4.1 Providing data to access a Repository Mirroring Tool server

By default, updates for SUSE Linux Enterprise Server are delivered by the SUSE Customer Center. If your network provides a Repository Mirroring Tool (RMT) server to provide a local update source, you need to equip the client with the server's URL. Client and server communicate solely via HTTPS protocol, therefore you also need to enter a path to the server's certificate if the certificate was not issued by a certificate authority.



Note: Non-interactive installation only

Providing parameters for accessing an RMT server is only needed for non-interactive installations. During an interactive installation the data can be provided during the installation (see [Section 9.7, "Registration"](#) for details).

regurl

URL of the RMT server. This URL has a fixed format of `https://FQN/center/regsvc/`. FQN needs to be a fully qualified host name of the RMT server. Example:

```
regurl=https://smt.example.com/center/regsvc/
```

Make sure the values you enter are correct. If regurl has not been specified correctly, the registration of the update source will fail.

regcert

Location of the RMT server's certificate. Specify one of the following locations:

URL

Remote location (HTTP, HTTPS or FTP) from which the certificate can be downloaded. In case regcert is not specified, it will default to `http://FQN/smt.crt` with FQN being the name of the RMT server. Example:

```
regcert=http://rmt.example.com/smt-ca.crt
```

local path

Absolute path to the certificate on the local machine. Example:

```
regcert=/data/inst/smt/smt-ca.cert
```

Interactive

Use ask to open a pop-up menu during the installation where you can specify the path to the certificate. Do not use this option with AutoYaST. Example

```
regcert=ask
```

Deactivate certificate installation

Use `done` if the certificate will be installed by an add-on product, or if you are using a certificate issued by an official certificate authority. For example:

```
regcert=done
```

8.4.2 Configuring an alternative data server for **supportconfig**

The data that supportconfig (see *Book "Administration Guide", Chapter 47 "Gathering system information for support"* for more information) gathers is sent to the SUSE Customer Center by default. It is also possible to set up a local server to collect this data. If such a server is available on your network, you need to set the server's URL on the client. This information needs to be entered at the boot prompt.

`supporturl`. URL of the server. The URL has the format `http://FQN/Path/`, where `FQN` is the fully qualified host name of the server and `Path` is the location on the server. For example:

```
supporturl=http://support.example.com/supportconfig/data/
```

8.4.3 Using IPv6 for the installation

By default you can only assign IPv4 network addresses to your machine. To enable IPv6 during installation, enter one of the following parameters at the boot prompt:

Accept IPv4 and IPv6

```
ipv6=1
```

Accept IPv6 only

```
ipv6only=1
```

8.4.4 Using a proxy for the installation

In networks enforcing the usage of a proxy server for accessing remote web sites, registration during installation is only possible when configuring a proxy server.

On systems with traditional BIOS, press **F4** on the boot screen and set the required parameters in the *HTTP Proxy* dialog.

On Systems with UEFI BIOS, provide the boot parameter proxy at the boot prompt:

1. On the boot screen, press **E** to edit the boot menu.
2. Append the proxy parameter to the linux line in the following format:

```
proxy=https://proxy.example.com:PORT
```

If the proxy server requires authentication, add the credentials as follows:

```
proxy=https://USER:PASSWORD@proxy.example.com:PORT
```

If the proxy server's SSL certificate cannot be verified, disable certificate checking with the sslcerts=0 boot parameter.

The outcome will be similar to the following:



FIGURE 8.3: GRUB OPTIONS EDITOR

3. Press **F10** to boot with the new proxy setting.

8.4.5 Enabling SELinux support

Enabling SELinux upon installation start-up enables you to configure it after the installation has been finished without having to reboot. Use the following parameters:

```
security=selinux selinux=1
```

8.4.6 Enabling the installer self-update

During installation and upgrade, YaST can update itself as described in [Section 9.2, “Installer self-update”](#) to solve potential bugs discovered after release. The `self_update` parameter can be used to modify the behavior of this feature.

To enable the installer self-update, set the parameter to `1`:

```
self_update=1
```

To use a user-defined repository, specify a URL:

```
self_update=https://updates.example.com/
```

8.4.7 Reusing LVM

As of SUSE Linux Enterprise 15 SP6, the installer no longer reuses pre-existing Logical Volume Manager (LVM) configurations in its *Guided Setup* for this can be confusing and lead to suboptimal setups. To reuse an existing LVM regardless, use the `YAST_REUSE_LVM` parameter or configure it manually in the *Expert Partitioner* ([Chapter 11, Expert Partitioner](#)).

8.4.8 Scale user interface for high DPI

If your screen uses a very high DPI, use the boot parameter `QT_AUTO_SCREEN_SCALE_FACTOR`. This scales font and user interface elements to the screen DPI.

```
QT_AUTO_SCREEN_SCALE_FACTOR=1
```

8.4.9 Using CPU mitigations

The boot parameter `mitigations` lets you control mitigation options for side-channel attacks on affected CPUs. Its possible values are:

`auto`. Enables all mitigations required for your CPU model, but does not protect against cross-CPU thread attacks. This setting may impact performance to some degree, depending on the workload.

`nosmt`. Provides the full set of available security mitigations. Enables all mitigations required for your CPU model. In addition, it disables Simultaneous Multithreading (SMT) to avoid side-

channel attacks across multiple CPU threads. This setting may further impact performance, depending on the workload.

`off`. Disables all mitigations. Side-channel attacks against your CPU are possible, depending on the CPU model. This setting has no impact on performance.

Each value comes with a set of specific parameters, depending on the CPU architecture, the kernel version, and on the vulnerabilities that need to be mitigated. Refer to the kernel documentation for details.

8.4.10 LUKS 2 Support

LUKS2 encryption is supported by the YaST installer as of SUSE Linux Enterprise 15 SP4, but needs to be enabled explicitly.

```
YAST_LUKS2_AVAILABLE
```

Alternatively, you can also enable LUKS2 in the YaST expert console. For more information, refer to [Section 11.2, “Device encryption”](#).

8.5 IBM Z

For IBM Z platforms, the system is booted (IPL, Initial Program Load) as described in [Section 5.3.4, “IPLing the SUSE Linux Enterprise Server installation system”](#). SUSE Linux Enterprise Server does not show a splash screen on these systems. During the installation, load the kernel, initrd, and parmfile manually. YaST starts with its installation screen when a connection has been established to the installation system via VNC, X, or SSH. Because there is no splash screen, kernel or boot parameters cannot be entered on screen, but must be specified in a parmfile (see [Section 5.5, “The parmfile—automating the system configuration”](#)).

`InstNetDev=osa`

Enter the type of interface to configure. Possible values are `osa`, `hsi`, `ctc`, `escon`, and `iucv` (CTC, ESCON, and IUCV are no longer supported).

For the interfaces of type `hsi` and `osa`, specify an appropriate netmask and an optional broadcast address:

```
Netmask=255.255.255.0
Broadcast=192.168.255.255
```


For the interfaces of type ctc, escon, and iucv (CTC, ESCON, and IUCV are no longer supported), enter the IP address of the peer:

```
Pointpoint=192.168.55.20
```

OsaInterface=<lcs|qdio>

For osa network devices, specify the host interface (qdio or lcs).

Layer2=<0|1>

For osa QDIO Ethernet and hsi devices, specify whether to enable (1) or disable (0) OSI Layer 2 support.

OSAHWAddr=02:00:65:00:01:09

For Layer 2-enabled osa QDIO Ethernet devices, either specify a MAC address manually or state OSAHWADDR= (with trailing white space) for the system default.

PortNo=<0|1>

For osa network devices, specify the port number (provided the device supports this feature). The default value is 0.

Each of the interfaces requires certain setup options:

- Interfaces ctc and escon (CTC and ESCON are no longer supported):

```
ReadChannel=0.0.0600
WriteChannel=0.0.0601
```

ReadChannel specifies the READ channel to use. WriteChannel specifies the WRITE channel.

- For the ctc interface (no longer supported), specify the protocol that should be used for this interface:

```
CTCProtocol=<0/1/2>
```

Valid entries would be:

<u>0</u>	Compatibility mode, also for non-Linux peers other than OS/390 and z/OS (this is the default mode)
<u>1</u>	Extended mode

- Network device type osa with interface lcs:

```
ReadChannel=0.0.0124
```

ReadChannel stands for the channel number used in this setup. A second port number can be derived from this by adding one to ReadChannel. Portnumber is used to specify the relative port.

- Interface iucv:

```
IUCVPeer=PEER
```

Enter the name of the peer machine.

- Network device type osa with interface qdio for OSA-Express Gigabit Ethernet:

```
ReadChannel=0.0.0700
WriteChannel=0.0.0701
DataChannel=0.0.0702
```

For ReadChannel, enter the number of the READ channel. For WriteChannel, enter the number of the WRITE channel. DataChannel specifies the DATA channel. Make sure that the READ channel carries an even device number.

- Interface hsi for HyperSockets and VM guest LANs:

```
ReadChannel=0.0.0800
WriteChannel=0.0.0801
DataChannel=0.0.0802
```

For ReadChannel, enter the appropriate number for the READ channel. For WriteChannel and DataChannel, enter the WRITE and DATA channel numbers.

8.6 More information

You can find more information about boot parameters in the openSUSE wiki at https://en.opensuse.org/SDB:Linuxrc#Parameter_Reference .

9 Installation steps

This chapter describes the procedure in which the data for SUSE Linux Enterprise Server is copied to the target device. Some basic configuration parameters for the newly installed system are set during the procedure. A graphical user interface will guide you through the installation. The procedure described in the following also applies to remote installation procedures as described in [Chapter 12, Remote installation](#). The text mode installation has the same steps and only looks different. For information about performing non-interactive automated installations, see *Book “AutoYaST Guide”*.

Before running the installer, read [Part I, “Installation preparation”](#). Depending on the architecture of your system, it describes the steps necessary to start the installation.

If you are a first-time user of SUSE Linux Enterprise Server, you should follow the default YaST proposals in most parts, but you can also adjust the settings as described here to fine-tune your system according to your preferences. Help for each installation step is provided by clicking *Help*.



Tip: Installation without a mouse

If the installer does not detect your mouse correctly, use **→|** for navigation, arrow keys to scroll, and **Enter** to confirm a selection. Various buttons or selection fields contain a letter with an underscore. Use **Alt + Letter** to select a button or a selection directly instead of navigating there with **→|**.

9.1 Overview

This section provides an overview of all installation steps. Each step contains a link to a more detailed description.

1. Before the installation starts, the installer may update itself. For details, see [Section 9.2, “Installer self-update”](#).
2. The actual installation starts with choosing the language and the product. For details, see [Section 9.3, “Language, keyboard and product selection”](#).
3. Accept the license agreement. For details, see [Section 9.4, “License agreement”](#).

4. IBM Z machines need to activate disks. For details, see [Section 9.5, “IBM Z: disk activation”](#).
5. Configure the network. This is only required when you need network access during the installation, and automatic network configuration via DHCP fails. If the automatic network configuration succeeds, this step is skipped. For details, see [Section 9.6, “Network settings”](#).
6. With a working network connection you can register the machine at the SUSE Customer Center or an RMT server. For details, see [Section 9.7, “Registration”](#).
7. Select the modules you want to enable for the machine. This impacts the availability of system roles in the next step and packages later on. For details, see [Section 9.8, “Extension and module selection”](#).
8. You can manually add repositories. For details, see [Section 9.9, “Add-on product”](#).
9. Select a role for your system. This defines the default list of packages to install and makes a suggestion for partitioning the hard disks. For details, see [Section 9.10, “System roles”](#).
10. Partition the hard disks of your system. For details, see [Section 9.11, “Partitioning”](#).
11. Choose a time zone. For details, see [Section 9.12, “Clock and time zone”](#).
12. Create a user. For details, see [Section 9.13, “Create new user”](#).
13. *(Optional)* Optionally, set a different password for the system administrator `root`. For details, see [Section 9.14, “Authentication for the system administrator root”](#).
14. In a final step, the installer presents an overview of all settings. If required, you can change them. For details, see [Section 9.15, “Installation settings”](#).
15. The installer copies all required data and informs you about the progress. For details, see [Section 9.16, “Performing the installation”](#).

9.2 Installer self-update


During the installation and upgrade process, YaST may update itself to solve bugs in the installer that were discovered after the release. This functionality is enabled by default; to disable it, set the boot parameter `self_update` to `0`. For more information, see [Section 8.4.6, “Enabling the installer self-update”](#).

Important: Quarterly media update: self-update disabled

The installer self-update is only available if you use the GM images of the Unified Installer and Packages ISOs. If you install from the ISOs published as quarterly update (they can be identified by the string QU in the name), the installer cannot update itself, because this feature is disabled in the update media.

Important: Networking during self-update

To download installer updates, YaST needs network access. By default, it tries to use DHCP on all network interfaces. If there is a DHCP server in the network, it will work automatically.

If you need a static IP setup, you can use the ifcfg boot argument. For more details, see the linuxrc documentation at <https://en.opensuse.org/Linuxrc> .

Tip: Language selection

The installer self-update runs before the language selection step. This means that progress and errors which happen during this process are displayed in English by default.

To use another language for this part of the installer, use the language boot parameter if available for your architecture, for example, language=de_DE. On machines equipped with a traditional BIOS, alternatively, press **F2** in the boot menu and select the language from the list.

Although this feature was designed to run without user intervention, it is worth knowing how it works. If you are not interested, you can jump directly to [Section 9.3, “Language, keyboard and product selection”](#) and skip the rest of this section.

9.2.1 Self-update process

The process can be broken down into two different parts:

1. Determine the update repository location.
2. Download and apply the updates to the installation system.

9.2.1.1 Determining the update repository location

Installer Self-Updates are distributed as regular RPM packages via a dedicated repository, so the first step is to find the repository URL.



Important: Installer self-update repository only

No matter which of the following options you use, only the installer self-update repository URL is expected, for example:

```
self_update=https://www.example.com/my_installer_updates/
```

Do not supply any other repository URL—for example the URL of the software update repository.

YaST will try the following sources of information:

1. The `self_update` boot parameter. (For more details, see [Section 8.4.6, “Enabling the installer self-update”](#).) If you specify a URL, it will take precedence over any other method.
2. The `/general/self_update_url` profile element in case you are using AutoYaST.
3. A registration server. YaST will query the registration server for the URL. The server to be used is determined in the following order:
 - a. By evaluating the `regurl` boot parameter ([Section 8.4.1, “Providing data to access a Repository Mirroring Tool server”](#)).
 - b. By evaluating the `/suse_register/reg_server` profile element if you are using AutoYaST.
 - c. By performing an SLP lookup. If an SLP server is found, YaST will ask you whether it should be used because there is no authentication involved and anybody on the local network can broadcast a registration server.
 - d. By querying the SUSE Customer Center.
4. If none of the previous attempts work, the fallback URL (defined in the installation media) will be used.

9.2.1.2 Downloading and applying the updates

When the update repository is determined, YaST checks whether an update is available. If it is, all the updates are downloaded and applied.

Finally, YaST restarts and displays the welcome screen. If no updates are available, the installation continues without restarting YaST.



Note: Update integrity

Update signatures will be checked to ensure integrity and authorship. If a signature is missing or invalid, you will be asked whether you want to apply the update.

9.2.1.3 Temporary self-update add-on repository

Some packages distributed in the self-update repository provide additional data for the installer, like installation defaults, system role definitions and similar. If the installer finds such packages in the self-update repository, a local temporary repository is created, to which those packages are copied. They are used during the installation. The temporary local repository is removed at the end of the installation. Its packages are *not* installed on the target system.

This additional repository is not displayed in the list of add-on products, but during installation it may still be visible as `SelfUpdate0` repository in the package management.

9.2.2 Custom self-update repositories

YaST can use a user-defined repository instead of the official repository by specifying a URL through the `self_update` boot parameter.

- HTTP/HTTPS and FTP repositories are supported.
- Starting with `yast2-installation-4.4.30`, the `reurl://` schema is supported, as a boot parameter or in an AutoYaST profile. The URL is relative to the main installation repository, and you may navigate the file tree with the usual `../` notation, for example `reurl:///../self_update`. This is useful when serving the packages via a local installation server, or when building a custom installation medium which includes a self-update repository.

The following examples assume the installation repository is at the medium root (/), and the self-update repository in the `self_update` subdirectory. This structure makes the `re-lurl://` portable, and it will work anywhere without changes as a boot parameter, copied to a USB stick, hard disk, network server, or in an AutoYaST profile.

Custom DVD/USB medium

Add the `self_update=relurl://self_update` boot option directly to the default boot parameters, and it will work properly even if the medium is copied to an USB stick, hard disk, or a network server.

Installation server

Assume that the installation packages are available via `http://example.com/repo` and a self-update repository is available at `http://example.com/self_update`. Then you can use the `http://example.com/repo` and `http://example.com/self_update` boot parameters, without having to change the `self_update` parameter when the repositories are moved to a different location.

- Only RPM-MD repositories are supported (required by RMT).
- Packages are not installed in the usual way: They are uncompressed only and no scripts are executed.
- No dependency checks are performed. Packages are installed in alphabetical order.
- Files from the packages override the files from the original installation media. This means that the update packages might not need to contain all files, only files that have changed. Unchanged files are omitted to save memory and download bandwidth.



Note: Only one repository

Currently, it is not possible to use more than one repository as source for installer self-updates.

9.3 Language, keyboard and product selection

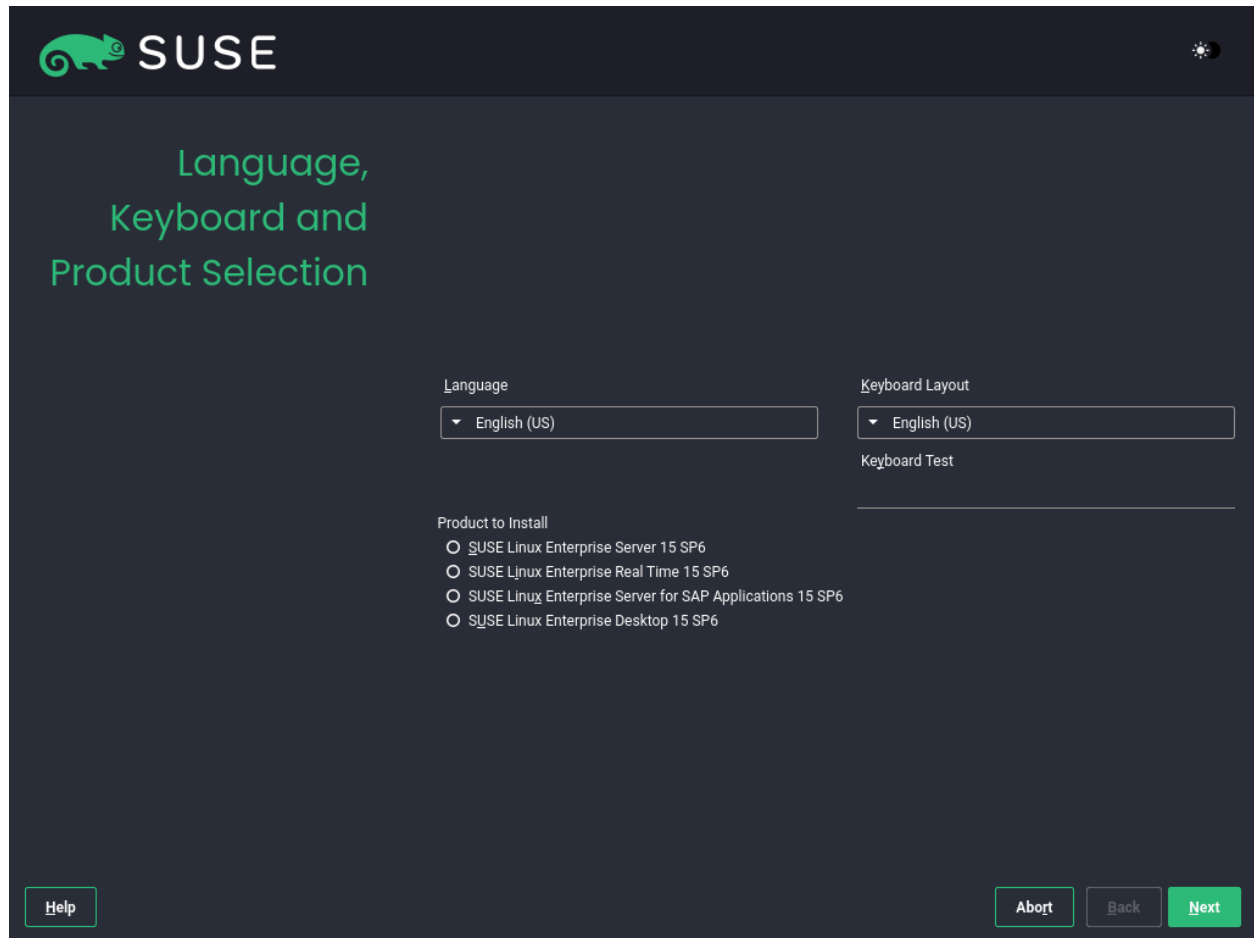


FIGURE 9.1: LANGUAGE, KEYBOARD AND PRODUCT SELECTION

The *Language* and *Keyboard Layout* settings are initialized with the language you chose on the boot screen. If you did not change the default, it will be English (US). Change the settings here, if necessary.

Changing the language automatically selects a corresponding keyboard layout. You can override this proposal by selecting a different keyboard layout from the drop-down box. Use the *Keyboard Test* text box to test the layout. The selected language also determines a time zone for the system clock. This setting can be modified later as described in *Book “Administration Guide”, Chapter 5 “Changing language and country settings with YaST”*.

With the Unified Installer, you can install all SUSE Linux Enterprise base products:


- SUSE Linux Enterprise Server 15 SP6 (covered here)
- SUSE Linux Enterprise Desktop 15 SP6 (for installation instructions, refer to <https://documentation.suse.com/sled/> )
- SUSE Linux Enterprise Real Time 15 SP6 (for installation instructions, refer to <https://documentation.suse.com/sle-rt/> )
- SUSE Linux Enterprise Server for SAP Applications 15 SP6 (for installation instructions, refer to <https://documentation.suse.com/sles-sap> )

Select a product for installation. You need to have a registration code for the respective product. In this document it is assumed you have chosen SUSE Linux Enterprise Server. Proceed with *Next*.



Tip: Light and high-contrast themes

If you have difficulties reading the labels in the installer, you can change the widget colors and theme.

Click the  button or press **Shift – F3** to open a theme selection dialog. Select a theme from the list and *Close* the dialog.

Shift – F4 switches to the color scheme for vision-impaired users. Press the buttons again to switch back to the default scheme.

9.4 License agreement

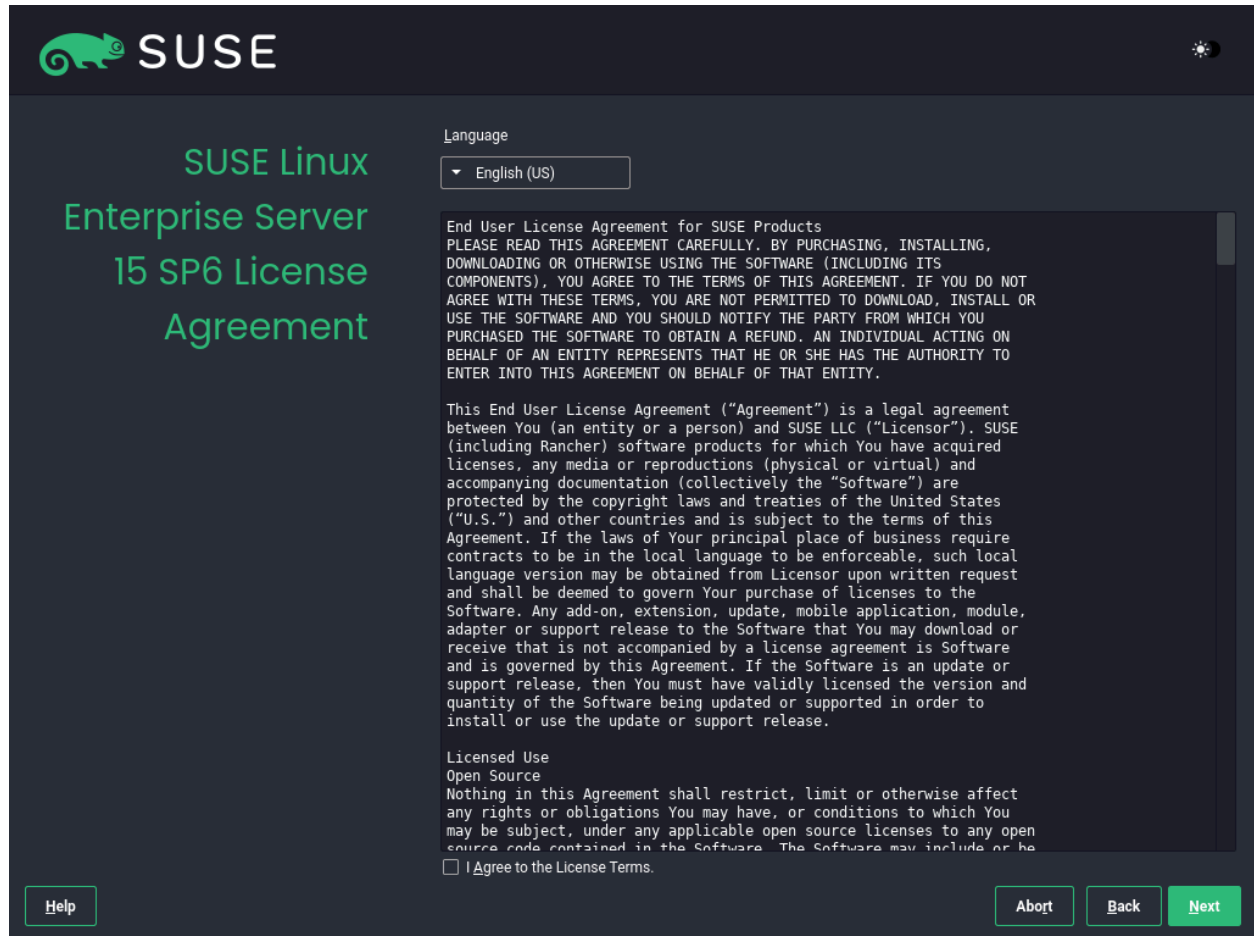


FIGURE 9.2: LICENSE AGREEMENT

Read the License Agreement. It is presented in the language you have chosen on the boot screen. Translations are available via the *License Language* drop-down box. If you agree to the terms, check *I Agree to the License Terms* and click *Next* to proceed with the installation. If you do not agree to the license agreement, you cannot install SUSE Linux Enterprise Server; click *Abort* to terminate the installation.

9.5 IBM Z: disk activation

When installing on IBM Z platforms, the language selection dialog is followed by a dialog to configure the attached hard disks.




FIGURE 9.3: DISK ACTIVATION

Select DASD, Fibre Channel Attached SCSI Disks (zFCP), or iSCSI for installation of SUSE Linux Enterprise Server. The DASD and zFCP configuration buttons are only available if the corresponding devices are attached. For instructions on how to configure iSCSI disks, refer to *Book “Storage Administration Guide”, Chapter 15 “Mass storage over IP networks: iSCSI”, Section 15.3 “Configuring iSCSI initiator”*.

You can also change the *Network Configuration* in this screen by launching the *Network Settings* dialog. Choose a network interface from the list and click *Edit* to change its settings. Use the tabs to configure DNS and routing. See *Book “Administration Guide”, Chapter 23 “Basic networking”, Section 23.4 “Configuring a network connection with YaST”* for more details.

9.5.1 Configuring DASD disks

Skip this step if you are not installing on IBM Z hardware.




FIGURE 9.4: DASD DISK MANAGEMENT

After selecting *Configure DASD Disks*, an overview lists all available DASDs. To get a clearer picture of the available devices, use the text box located above the list to specify a range of channels to display. To filter the list according to such a range, select *Filter*.

Specify the DASDs to use for the installation by selecting the corresponding entries in the list. Use *Select All* to select all DASDs currently displayed. Activate and make the selected DASDs available for the installation by selecting *Perform Action* › *Activate*. To format the DASDs, select *Perform Action* › *Format*. Alternatively, use the YaST partitioner later as described in [Section 11.1, “Using the Expert Partitioner”](#).

9.5.2 Configuring zFCP disks

Skip this step if you are not installing on IBM Z hardware.

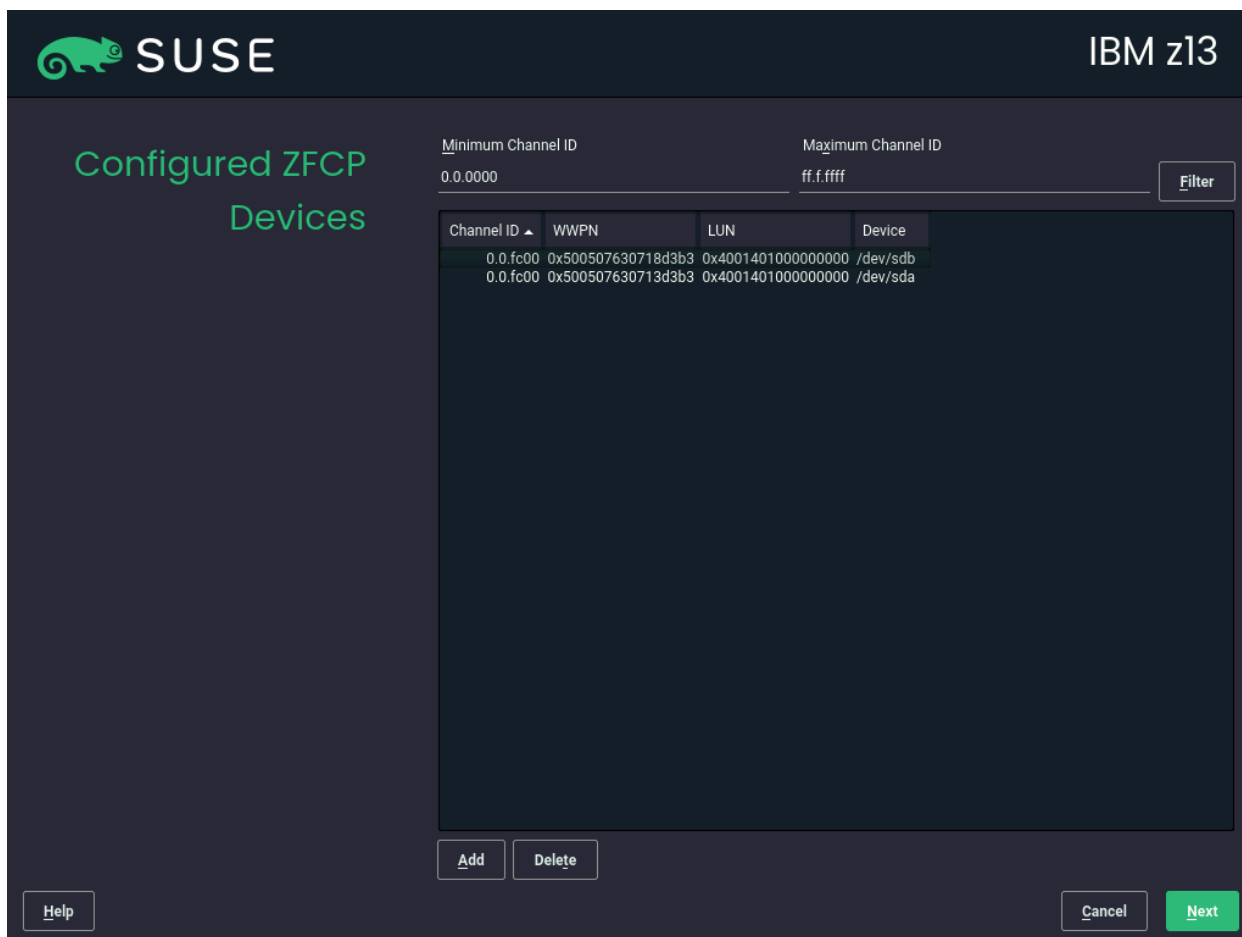


FIGURE 9.5: CONFIGURED ZFCP DEVICES

After selecting *Configure zFCP Disks*, a dialog with a list of the zFCP disks available on the system opens. In this dialog, select *Add* to open another dialog in which to enter zFCP parameters.

To make a zFCP disk available for the SUSE Linux Enterprise Server installation, choose an available *Channel Number* from the drop-down box. *Get WWPNs* (World Wide Port Number) and *Get LUNs* (Logical Unit Number) return lists with available WWPNs and FCP-LUNs, respectively, to choose from. Automatic LUN scanning only works with NPIV enabled.

When completed, exit the zFCP dialog with *Next* and the general hard disk configuration dialog with *Finish* to continue with the rest of the configuration.

9.6 Network settings

After booting into the installation, the installation routine is set up. During this setup, an attempt to configure at least one network interface with DHCP is made. In case this attempt has failed, the *Network Settings* dialog launches now.




FIGURE 9.6: NETWORK SETTINGS

Choose a network interface from the list and click *Edit* to change its settings. Use the tabs to configure DNS and routing. See *Book "Administration Guide", Chapter 23 "Basic networking", Section 23.4 "Configuring a network connection with YaST"* for more details. On IBM Z this dialog does not start automatically. It can be started in the *Disk Activation* step.

In case DHCP was successfully configured during installation setup, you can also access this dialog by clicking *Network Configuration* at the *SUSE Customer Center Registration* and the *Installation Settings* step. It lets you change the automatically provided settings.



Note: Network configuration with boot parameters

If at least one network interface has been configured via boot parameters (see [Section 8.3.2, “Configuring the network interface”](#)), automatic DHCP configuration is disabled and the boot parameter configuration is imported and used.



Tip: Accessing network storage or local RAID

To access a SAN or a local RAID during the installation, you can use the `libstorage` command line client for this purpose:

1. Switch to a console with `Ctrl – Alt – F2`.
2. Install the `libstoragemgmt` extension by running `extend libstoragemgmt`.
3. Now you have access to the `lsmcli` command. For more information, run `lsmcli --help`.
4. To return to the installer, press `Alt – F7`.

Supported are Netapp Ontap, all SMI-S compatible SAN providers, and LSI MegaRAID.

9.7 Registration

To get technical support and product updates, you need to register and activate SUSE Linux Enterprise Server with the SUSE Customer Center or a local registration server. Registering your product at this stage also grants you immediate access to the update repository. This enables you to install the system with the latest updates and patches available.

When registering, repositories and dependencies for modules and extensions are loaded from the registration server.

From this dialog, you can switch to the YaST *Network Settings* module by clicking *Network Configuration*. For details, see *Book “Administration Guide”, Chapter 23 “Basic networking”, Section 23.4 “Configuring a network connection with YaST”*.

If you are offline or want to skip registration, activate *Skip Registration*. See [Section 9.7.3, “Installing without registration”](#) for instructions.

9.7.1 Manual registration

To register with the SUSE Customer Center, provide the *E-mail Address* associated with your SCC account and the *Registration Code* for SUSE Linux Enterprise Server.

If your organization offers a local registration server, you may register there. Activate *Register System via local SMT Server* and either choose a URL from the drop-down box or type in an address. Proceed with *Next*.

To register with the SUSE Customer Center, enter your *Registration Code* for SUSE Linux Enterprise Server. If your organization provides a local registration server, you may register there. Activate *Register System via local RMT Server* and either choose a URL from the drop-down box or type in an address.

Start the registration process with *Next*.

The screenshot displays the SUSE registration interface. At the top left is the SUSE logo. The main heading is 'Registration'. In the top right corner, there is a button labeled 'Network Configuration...'. The central text reads 'SUSE Linux Enterprise Server 15 SP6' followed by 'Please select your preferred method of registration.' There are three radio button options: 1. 'Register System via scc.suse.com' (selected): Includes input fields for 'E-mail Address' and 'Registration Code'. 2. 'Register System via local RMT Server': Includes a dropdown menu for 'Local Registration Server URL' with the value 'https://rmt.example.com'. 3. 'Skip Registration'. At the bottom, there are four buttons: 'Help', 'Abort', 'Back', and 'Next'.

FIGURE 9.7: SUSE CUSTOMER CENTER REGISTRATION



Tip: Installing product patches at installation time

After SUSE Linux Enterprise Server has been successfully registered, you are asked whether to install the latest available online updates during the installation. If you choose *Yes*, the system will be installed with the most current packages without having to apply updates after installation. It is recommended to enable this option.



Note: Firewall settings for receiving updates

By default, the firewall on SUSE Linux Enterprise Server only blocks incoming connections. If your system is behind another firewall that blocks outgoing traffic, make sure to allow connections to <https://scc.suse.com/> and <https://updates.suse.com> on ports 80 and 443 in order to receive updates.

If the system is successfully registered during installation, YaST disables repositories from local installation media such as CD/DVD or flash disks when the installation completes. This prevents problems caused by the missing installation source and ensures that you always get the latest updates from the online repositories.

9.7.2 Loading registration codes from USB storage

To make the registration more convenient, you can also store your registration codes on a USB storage device such as a flash disk. YaST will automatically pre-fill the corresponding text box. This is particularly useful when testing the installation or if you need to register many systems or extensions.

Create a file named `regcodes.txt` or `regcodes.xml` on the USB disk. If both are present, the XML takes precedence.

In that file, identify the product with the name returned by `zypper search --type product` and assign it a registration code as follows:

EXAMPLE 9.1: `regcodes.txt`

```
SLES    cc36aae1
SLED    309105d4

sle-we  5eedd26a
sle-live-patching 8c541494
```

EXAMPLE 9.2: regcodes.xml

```
<?xml version="1.0"?>
<profile xmlns="http://www.suse.com/1.0/yast2ns"
  xmlns:config="http://www.suse.com/1.0/configns">
  <suse_register>
    <addons config:type="list">
      <addon>
<name>SLES</name>
<reg_code>cc36aae1</reg_code>
      </addon>
      <addon>
<name>SLED</name>
<reg_code>309105d4</reg_code>
      </addon>
      <addon>
<name>sle-we</name>
<reg_code>5eedd26a</reg_code>
      </addon>
      <addon>
<name>sle-live-patching</name>
<reg_code>8c541494</reg_code>
      </addon>
    </addons>
  </suse_register>
</profile>
```

Note that SLES and SLED are not extensions, but listing them as add-ons allows for combining several base product registration codes in a single file. See *Book "AutoYaST Guide", Chapter 4 "Configuration and installation options", Section 4.3.1 "Extensions"* for details.



Note: Limitations

Currently flash disks are only scanned during installation or upgrade, but not when registering a running system.

9.7.3 Installing without registration

If you are offline or want to skip registration, activate *Skip Registration*. Accept the warning with *OK* and proceed with *Next*.

! Important: Skipping the registration

Your system and extensions need to be registered to retrieve updates and to be eligible for support. Skipping the registration is only possible when installing from the SLE-15-SP6-Full-ARCH-GM-media1.iso image.




FIGURE 9.8: INSTALLING WITHOUT REGISTRATION

📎 Note: Registering SUSE Linux Enterprise Server

Your system and extensions need to be registered to retrieve updates and to be eligible for support. If you do not register during the installation, you can do so at any time later from the running system. To do so, run *YaST > Product Registration*.



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
```

IMAGE needs to be replaced with the path to the [SLE-15-SP6-Online-ARCH-GM-media1.iso](#) or [SLE-15-SP6-Full-ARCH-GM-media1.iso](#) image file. *FLASH_DISK* needs to be replaced with the flash device. To identify the device, insert it and run:

```
# grep -Ff <(hwinfo --disk --short) <(hwinfo --usb --short)
disk:
/dev/sdc          General USB Flash Disk
```

Make sure the size of the device is sufficient for the desired image. You can check the size of the device with:

```
# fdisk -l /dev/sdc | grep -e "^/dev"
/dev/sdc1 *      2048 31490047 31488000  15G 83 Linux
```

In this example, the device has a capacity of 15 GB. The command to use for the [SLE-15-SP6-Full-ARCH-GM-media1.iso](#) would be:

```
dd if=SLE-15-SP6-Full-ARCH-GM-media1.iso of=/dev/sdc bs=4M && sync
```

The device must not be mounted when running the **dd** command. Note that all data on the partition will be erased!

9.8 Extension and module selection

In this dialog the installer lists modules and extensions that are available for SUSE Linux Enterprise Server. Modules are components that allow you to customize the product according to your needs. They are included in your SUSE Linux Enterprise Server subscription. Extensions add functionality to your product. They must be purchased separately.

The availability of certain modules or extensions depends on the product you chose in the first step of this installation. For a description of the modules and their lifecycles, select a module to see the accompanying text. More detailed information is available in the [Modules and Extensions Quick Start \(https://documentation.suse.com/sles-15/html/SLES-all/article-modules.html\)](https://documentation.suse.com/sles-15/html/SLES-all/article-modules.html).

The selection of modules indirectly affects the scope of the installation, because it defines which software sources (repositories) are available for installation and in the running system.

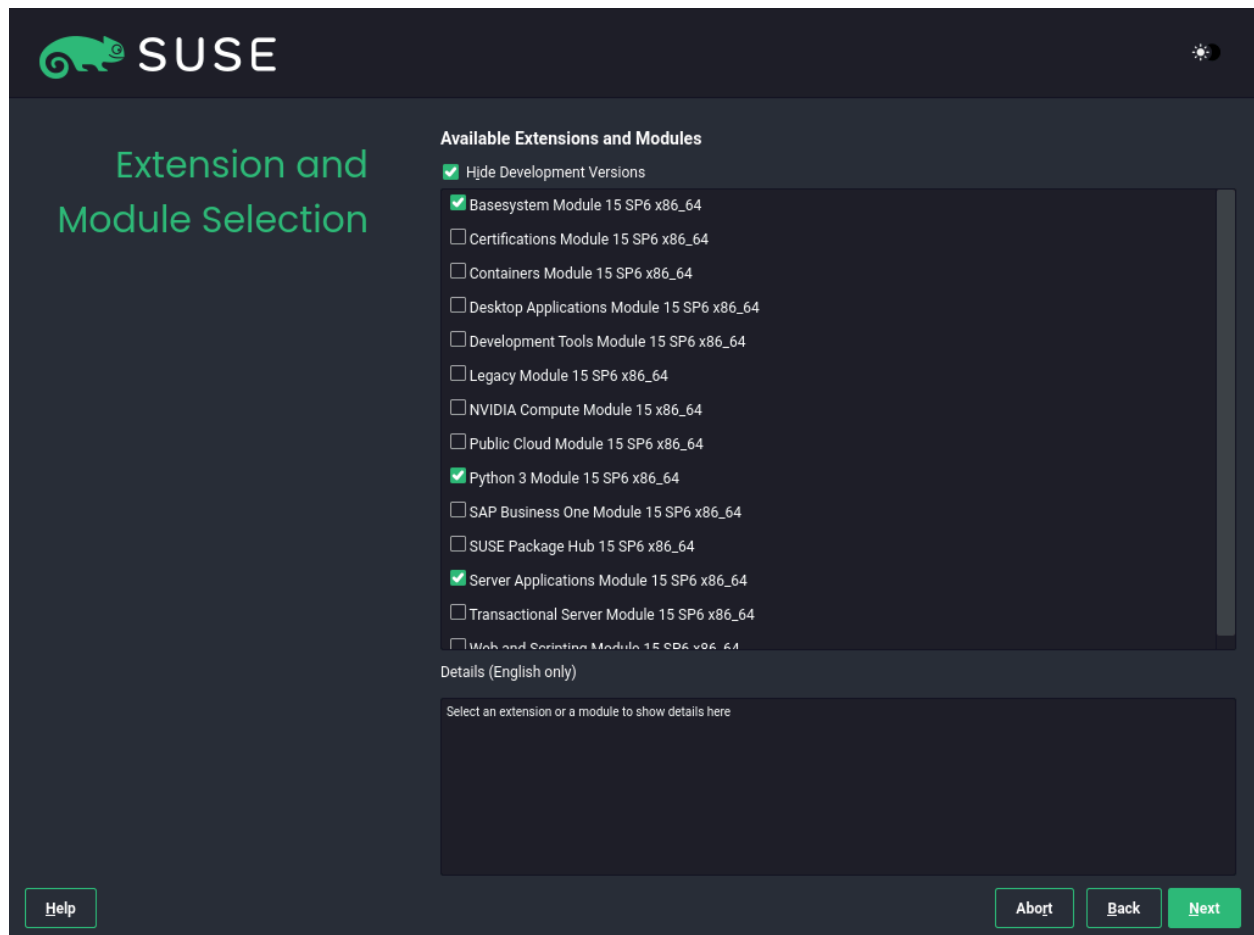


FIGURE 9.9: EXTENSION AND MODULE SELECTION

The following modules and extensions are available for SUSE Linux Enterprise Server:

Basesystem Module

This module adds a basic system on top of the Unified Installer. It is required by all other modules and extensions. The scope of an installation that only contains the base system is comparable to the installation pattern *minimal system* of previous SUSE Linux Enterprise Server versions. This module is selected for installation by default and should not be de-selected.

Dependencies: None

Certifications Module

Contains the FIPS certification packages.

Dependencies: Server Applications

Confidential Computing Technical Preview

Contains packages related to confidential computing.

Dependencies: Basesystem

Containers Module

Contains support and tools for containers.

Dependencies: Basesystem

Desktop Applications Module

Adds a graphical user interface and essential desktop applications to the system.

Dependencies: Basesystem

Development Tools Module

Contains compilers (including `gcc`) and libraries required for compiling and debugging applications. Replaces the former Software Development Kit (SDK).

Dependencies: Basesystem, Desktop Applications

High Performance Computing (HPC) Module

Provides specific tools commonly used for high performance, numerically intensive workloads.

Dependencies: Basesystem

Legacy Module


Helps you with migrating applications from earlier versions of SUSE Linux Enterprise Server and other systems to SLES 15 SP6, by providing packages which are discontinued on SUSE Linux Enterprise. Packages in this module are selected based on the requirement for migration and the level of complexity of configuration.

This module is recommended when migrating from a previous product version.

Dependencies: Basesystem, Server Applications

NVIDIA Compute Module

Contains the NVIDIA CUDA (Compute Unified Device Architecture) drivers.

The software in this module is provided by NVIDIA under the [CUDA End User License Agreement](http://docs.nvidia.com/cuda/eula/) (<http://docs.nvidia.com/cuda/eula/>)  and is not supported by SUSE.

Dependencies: Basesystem

Public Cloud Module

Contains all tools required to create images for deploying SUSE Linux Enterprise Server in cloud environments such as Amazon Web Services (AWS), Microsoft Azure, Google Compute Platform, or OpenStack.

Dependencies: Basesystem, Server Applications

Python 3 Module

This module contains the most recent version of the selected Python 3 packages.

Dependencies: Basesystem

SAP Business One Server

This module contains packages and system configuration specific to SAP Business One Server. It is maintained and supported by the SUSE Linux Enterprise Server product subscription.

Dependencies: Basesystem, Server Applications, Desktop Applications, Development Tools

Server Applications Module

Adds server functionality by providing network services such as DHCP server, name server, or Web server. This module is selected for installation by default; deselecting it is not recommended.

Dependencies: Basesystem

SUSE Linux Enterprise High Availability

Adds clustering support for mission critical setups to SUSE Linux Enterprise Server. This extension requires a separate license key.

Dependencies: Basesystem, Server Applications

SUSE Linux Enterprise Live Patching

Adds support for performing critical patching without having to shut down the system. This extension requires a separate license key.

Dependencies: Basesystem, Server Applications

SUSE Linux Enterprise Workstation Extension

Extends the functionality of SUSE Linux Enterprise Server with packages from SUSE Linux Enterprise Desktop, like additional desktop applications (office suite, e-mail client, graphical editor, etc.) and libraries. It allows to combine both products to create a fully featured workstation. This extension requires a separate license key.

Dependencies: Basesystem, Desktop Applications

SUSE Package Hub

Provides access to packages for SUSE Linux Enterprise Server maintained by the openSUSE community. These packages are delivered without L3 support and do not interfere with the supportability of SUSE Linux Enterprise Server. For more information, refer to <https://packagehub.suse.com/>.

Dependencies: Basesystem

Transactional Server Module

Adds support for transactional updates. Updates are either applied to the system all together in a single transaction, or not. This happens without influencing the running system. If an update fails, or if the successful update is deemed to be incompatible or otherwise incorrect, it can be discarded to immediately return the system to its previous functioning state.

Dependencies: Basesystem

Web and Scripting Module

Contains packages intended for a running Web server.

Dependencies: Basesystem, Server Applications

Some modules depend on the installation of other modules. Therefore, when selecting a module, other modules may be selected automatically to fulfill dependencies.

Depending on the product, the registration server can mark modules and extensions as recommended. Recommended modules and extensions are preselected for registration and installation. To avoid installing these recommendations, deselect them manually.

Select the modules and extensions you want to install and proceed with *Next*. In case you have chosen one or more extensions, you will be prompted to provide the respective registration codes. Depending on your choice, it may also be necessary to accept additional license agreements.



Important: Default modules for offline installation

When performing an offline installation from the SLE-15-SP6-Full-ARCH-GM-media1.iso, only the *Basesystem Module* is selected by default. To install the complete default package set of SUSE Linux Enterprise Server, additionally select the *Server Applications Module*.

9.9 Add-on product

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

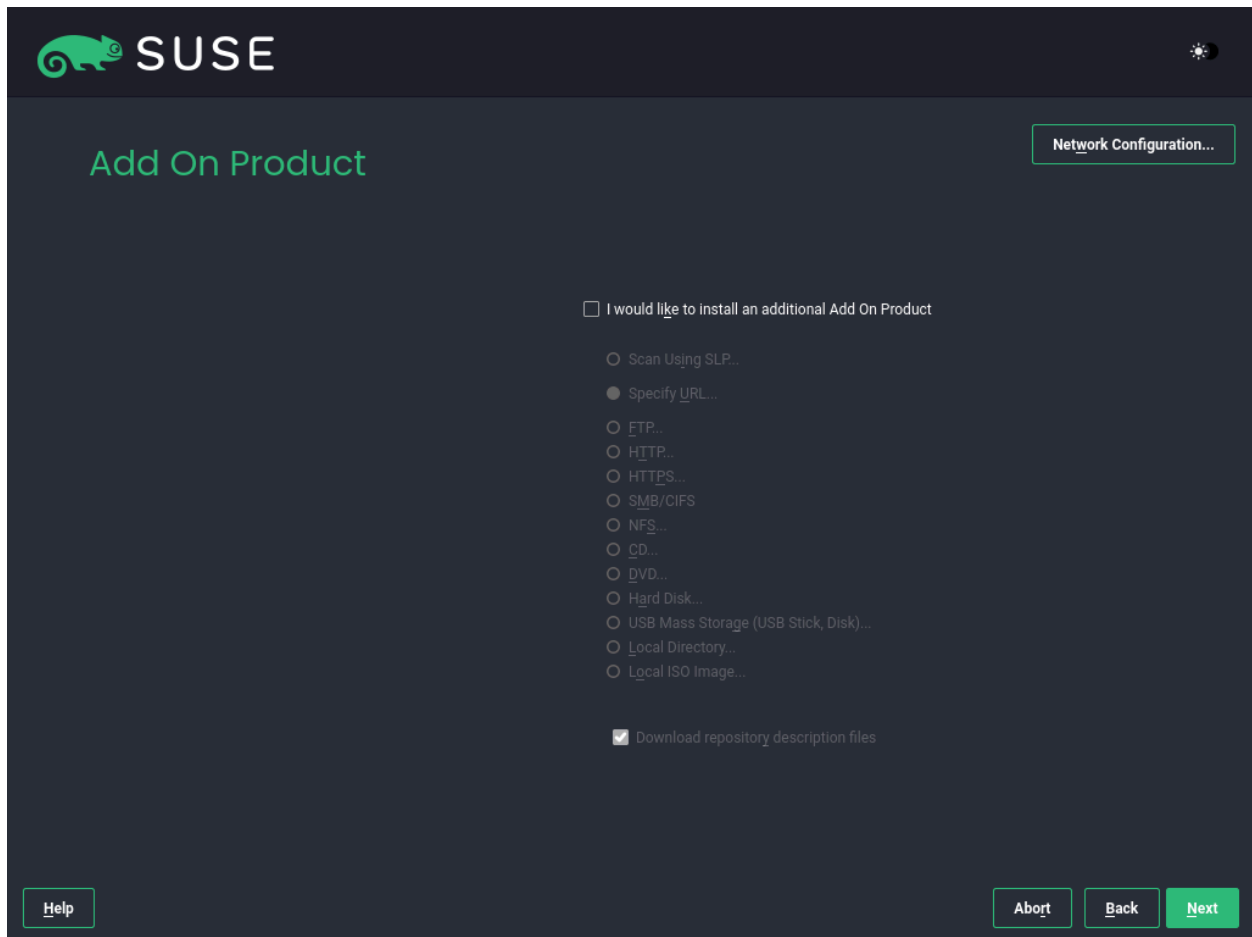



FIGURE 9.10: ADD-ON PRODUCT

From this dialog, you can switch to the YaST *Network Settings* module by clicking *Network Configuration*. For details, see Book “Administration Guide”, Chapter 23 “Basic networking”, Section 23.4 “Configuring a network connection with YaST”.



Tip: Adding drivers during the installation

You can also add driver update repositories via the *Add On Product* dialog. Driver updates for SUSE Linux Enterprise are provided at <https://drivers.suse.com/> . These drivers have been created via the SUSE SolidDriver Program.

If you do not want to install add-ons, proceed with *Next*. Otherwise activate *I would like to install an additional Add On Product*. Specify the Media Type by choosing from CD, DVD, Hard Disk, USB Mass Storage, a Local Directory or a Local ISO Image. If network access has been configured you can choose from additional remote sources such as HTTP, SLP, FTP, etc. Alternatively you may directly specify a URL. Check *Download repository description files* to download the files describing the repository now. If deactivated, they will be downloaded after the installation starts. Proceed with *Next* and insert a CD or DVD if required.

Depending on the add-on's content, it may be necessary to accept additional license agreements.

9.10 System roles

To simplify the installation, the installer offers predefined use cases that tailor the system for the selected scenario.




FIGURE 9.11: **SYSTEM ROLE**

Choose the *System Role* that meets your requirements best. The availability of system roles depends on your selection of modules and extensions. The dialog is omitted under the following conditions:

- The combination of base product and modules does not allow roles to be chosen.
- The combination of base product and modules only allows a single role.

With the default selection, the following system roles are available:

Text Mode

This option installs a basic SLES without a desktop environment but with a rich set of command line tools.

Dependencies: Basesystem

Minimal

Select this role if you want a very small installation with only the basic command line tools.

Dependencies: None

KVM Virtualization Host

Select this scenario when installing on a machine that should serve as a KVM host that can run other virtual machines. /var/lib/libvirt will be placed on a separate partition and the firewall and Kdump will be disabled.

Dependencies: Basesystem, Server Applications

Xen Virtualization Host

Select this scenario when installing on a machine that should serve as a Xen host that can run other virtual machines. /var/lib/libvirt will be placed on a separate partition and the firewall and Kdump will be disabled.

Dependencies: Basesystem, Server Applications

9.11 Partitioning

9.11.1 Important information



Warning: Read this section carefully

Read this section carefully before continuing with [Section 9.11.2, “Suggested partitioning”](#).

Custom partitioning on UEFI machines

A UEFI machine *requires* an EFI system partition that must be mounted to /boot/efi. This partition must be formatted with the FAT32 file system.

If an EFI system partition is already present on your system (for example from a previous Windows installation) use it by mounting it to /boot/efi without formatting it.

If no EFI system partition is present on your UEFI machine, make sure to create it. The EFI system partition must be a physical partition or RAID 1. Other RAID levels, LVM and other technologies are not supported. It needs to be formatted with the FAT32 file system.

Custom partitioning and Snapper

If the root partition is larger than 16 GB, SUSE Linux Enterprise Server by default enables file system snapshots.

SUSE Linux Enterprise Server uses Snapper together with Btrfs for this feature. Btrfs needs to be set up with snapshots enabled for the root partition.

If the disk is smaller than 16 GB, all Snapper features and automatic snapshots are disabled to prevent the system partition `/` from running out of space.

Being able to create system snapshots that enable rollbacks requires important system directories to be mounted on a single partition, for example `/usr` and `/var`. Only directories that are excluded from snapshots may reside on separate partitions, for example `/usr/local`, `/var/log`, and `/tmp`.

If snapshots are enabled, the installer will automatically create single snapshots during and immediately after the installation.

For details, see Book “Administration Guide”, Chapter 10 “System recovery and snapshot management with Snapper”.



Important: Btrfs snapshots and root partition size

Snapshots may take considerable storage space. Generally, the older a snapshot is or the larger the changeset it covers, the more storage space the snapshot takes. And the more snapshots you keep, the more disk space you need.

To prevent the root partition running full with snapshot data, you need to make sure it is big enough. In case you do frequent updates or other installations, consider at least 30 GB for the root partition. If you plan to keep snapshots activated for a system upgrade or a service pack migration (to be able to roll back), you should consider 40 GB or more.

Btrfs data volumes

Using Btrfs for data volumes is supported on SUSE Linux Enterprise Server 15 SP6. For applications that require Btrfs as a data volume, consider creating a separate file system with quota groups disabled. This is already the default for non-root file systems.

Btrfs on an encrypted root partition

The default partitioning setup suggests the root partition as Btrfs. To encrypt the root partition, make sure to use the GPT partition table type instead of the MSDOS type. Otherwise the GRUB2 boot loader may not have enough space for the second stage loader.

IBM Z: Using minidisks in z/VM

If SUSE Linux Enterprise Server is installed on minidisks in z/VM, which reside on the same physical disk, the access path of the minidisks (`/dev/disk/by-id/`) is not unique. This is because it represents the ID of the physical disk. If two or more minidisks are on the same physical disk, they all have the same ID.

To avoid problems when mounting minidisks, always mount them either *by path* or *by UUID*.

IBM Z: Using FBA DASDs in z/VM

If SUSE Linux Enterprise Server is installed on FBA DASDs in z/VM, a suggested partitioning cannot be provided. Instead, choose *Expert Partitioner* › *Start with existing partitions*.

FBA DASD comes with an implicit partition that must not be deleted, but should be reused without any changes. Do *not* repartition the FBA DASD.

IBM Z: LVM root file system

If you configure the system with a root file system on LVM or software RAID array, you must place `/boot` on a separate, non-LVM or non-RAID partition, otherwise the system will fail to boot. The recommended size for such a partition is 500 MB and the recommended file system is Ext4.


IBM POWER: Installing on systems with multiple Fibre Channel disks

If more than one disk is available, the partitioning scheme suggested during the installation puts the PReP and BOOT partitions on different disks. If these disks are Fibre Channel disks, the GRUB boot loader is not able to find the BOOT partition and the system cannot be booted.

When prompted to select the partition scheme during the installation, choose *Guided Setup* and verify that only one disk is selected for installation. Alternatively, run the *Expert Partitioner* and manually set up a partitioning scheme that has PReP and BOOT on a single disk.

Supported software RAID volumes

Installing to and booting from existing software RAID volumes is supported for Disk Data Format (DDF) volumes and Intel Matrix Storage Manager (IMSM) volumes. IMSM is also known by the following names:

- Intel Rapid Storage Technology
- Intel Matrix Storage Technology
- Intel Application Accelerator / Intel Application Accelerator RAID Edition
- Intel Virtual RAID on CPU (Intel VROC, see <https://www.intel.com/content/www/us/en/support/articles/000024498/memory-and-storage/ssd-software.html>  for more details)

Mount points for FCoE and iSCSI devices

FCoE and iSCSI devices will appear asynchronously during the boot process. While the `initrd` guarantees that those devices are set up correctly for the root file system, there are no such guarantees for any other file systems or mount points like `/usr`. Hence any system mount points like `/usr` or `/var` are not supported. To use those devices, ensure correct synchronization of the respective services and devices.

9.11.2 Suggested partitioning

Define a partition setup for SUSE Linux Enterprise Server in this step.

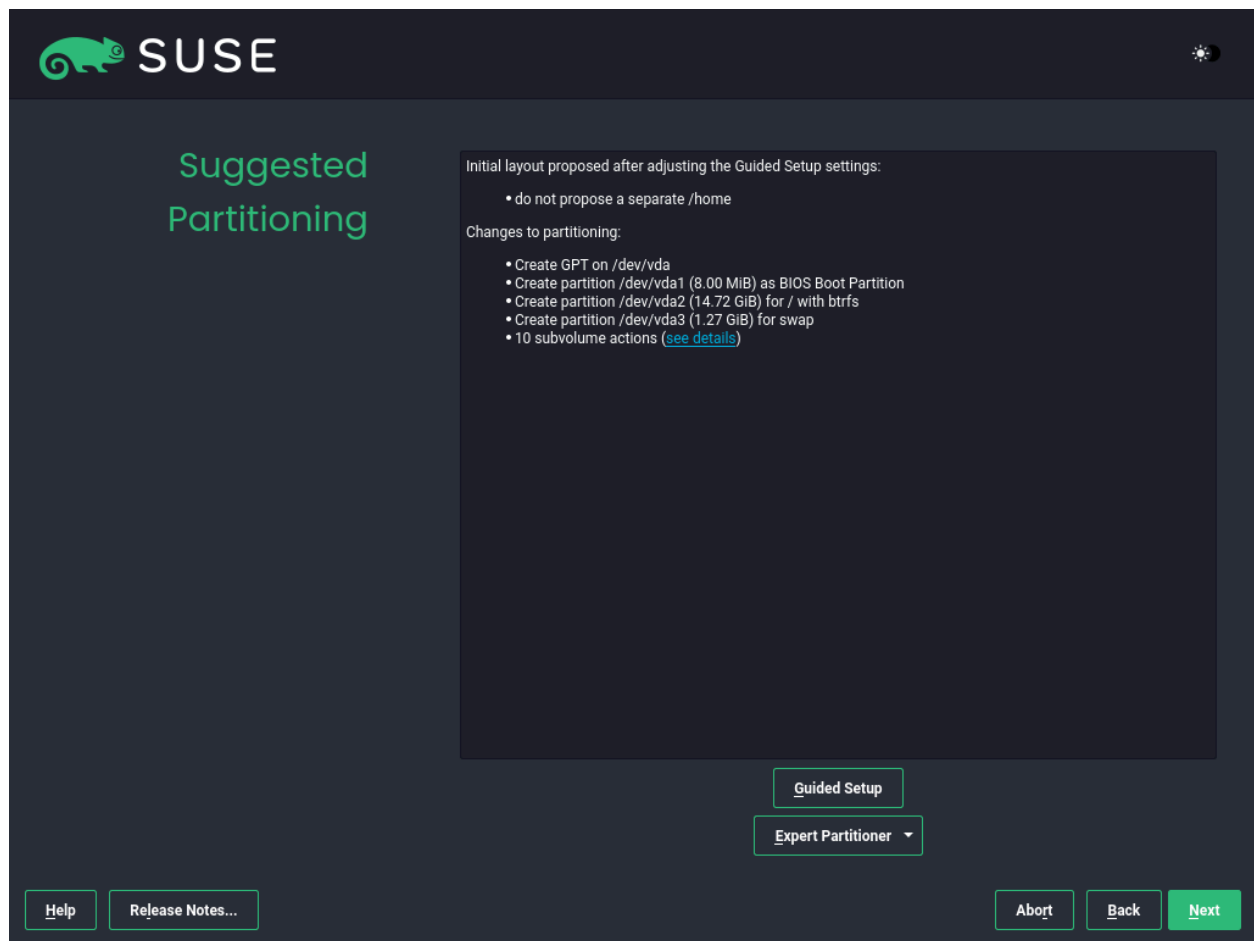


FIGURE 9.12: SUGGESTED PARTITIONING

Depending on the system role, the installer creates a proposal for one of the disks available. All proposals contain a root partition formatted with Btrfs (with snapshots enabled) and a swap partition. The GNOME desktop and the text mode proposals create a separate home partition on disks larger than 20 GB. The system roles for virtualization hosts create a separate partition

for `/var/lib/libvirt`, the directory that hosts the image files by default. If one or more swap partitions have been detected on the available hard disks, these existing ones will be used (rather than proposing a new swap partition). You have several options to proceed:

Next

To accept the proposal without any changes, click *Next* to proceed with the installation workflow.

Guided setup

To adjust the proposal, choose *Guided Setup*. First, choose which hard disks and partitions to use. In the *Partitioning Scheme* screen, you can enable Logical Volume Management (LVM) and activate disk encryption. Afterward specify the *Filesystem Options*. You can adjust the file system for the root partition and create a separate home and swap partitions. If you plan to suspend your machine, make sure to create a separate swap partition and check *Enlarge to RAM Size for Suspend*. If the root file system format is Btrfs, you can also enable or disable Btrfs snapshots here.

Expert Partitioner

To create a custom partition setup click *Expert Partitioner*. Select either *Start with Current Proposal* if you want start with the suggested disk layout, or *Start with Existing Partitions* to ignore the suggested layout and start with the existing layout on the disk. You can *Add*, *Edit*, *Resize*, or *Delete* partitions.

You can also set up logical volume management (LVM), configure software RAID and device mapping (DM), encrypt partitions, mount NFS shares and manage tmpfs volumes with the *Expert Partitioner*. To fine-tune settings such as the subvolume and snapshot handling for each Btrfs partition, choose *Btrfs*. For more information about custom partitioning and configuring advanced features, refer to [Section 11.1, “Using the Expert Partitioner”](#).



Warning: Disk space units

Note that for partitioning purposes, disk space is measured in binary units, rather than in decimal units. For example, if you enter sizes of 1GB, 1GiB or 1G, they all signify 1 GiB (Gibibyte), as opposed to 1 GB (Gigabyte).

Binary

1 GiB = 1 073 741 824 bytes.

Decimal

1 GB = 1 000 000 000 bytes.

Difference

1 GiB \approx 1.07 GB.

9.12 Clock and time zone

In this dialog, select your region and time zone. Both are preselected according to the installation language.

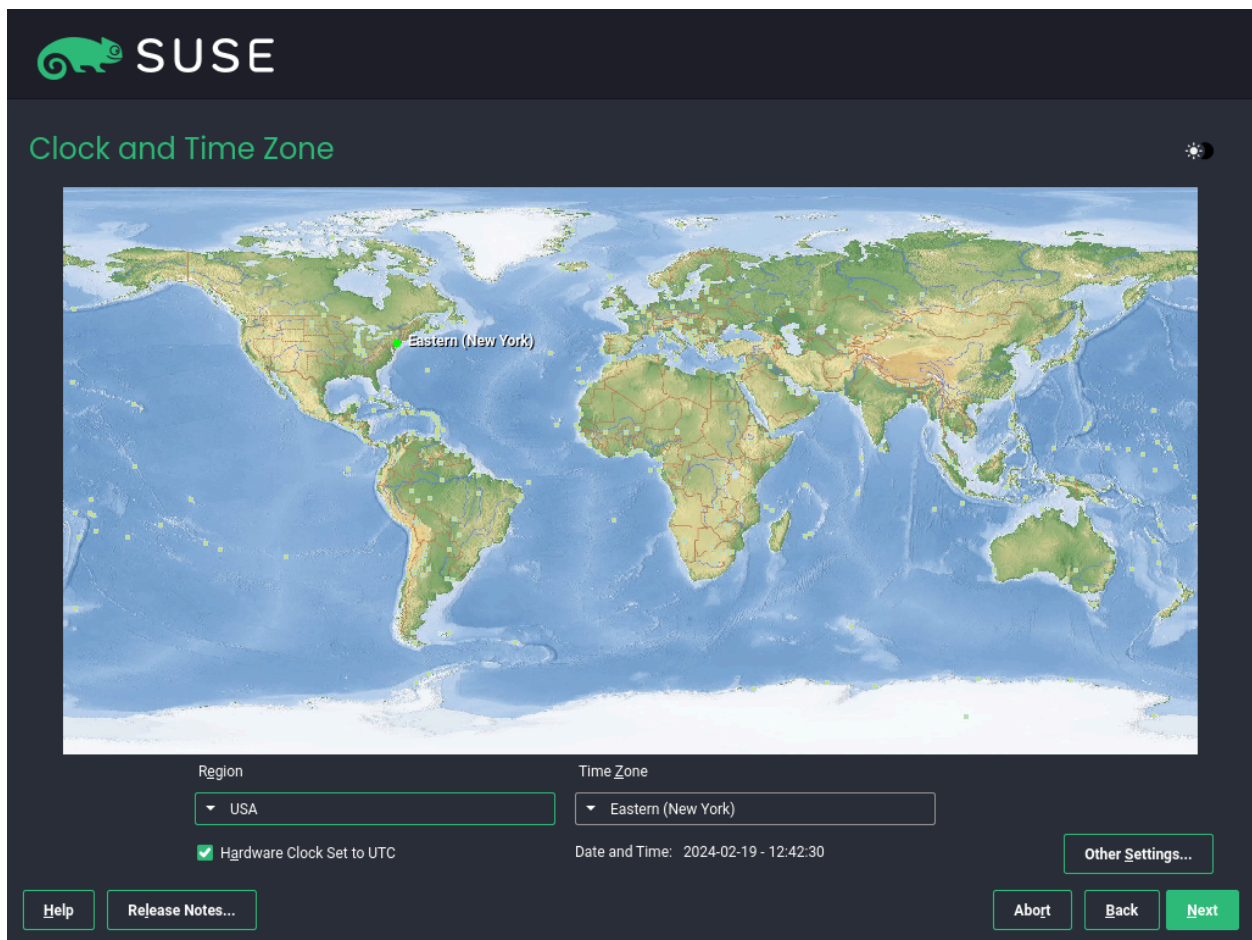


FIGURE 9.13: CLOCK AND TIME ZONE

To change the preselected values, either use the map or the drop-down boxes for *Region* and *Time Zone*. When using the map, point the cursor at the rough direction of your region and left-click to zoom. Now choose your country or region by left-clicking. Right-click to return to the world map.

To set up the clock, choose whether the *Hardware Clock is Set to UTC*. If you run another operating system on your machine, such as Microsoft Windows, it is likely your system uses local time instead. If you run Linux on your machine, set the hardware clock to UTC and have the switch from standard time to daylight saving time performed automatically.

Important: Set the hardware clock to UTC

The switch from standard time to daylight saving time (and vice versa) can only be performed automatically when the hardware clock (CMOS clock) is set to UTC. This also applies if you use automatic time synchronization with NTP, because automatic synchronization will only be performed if the time difference between the hardware and system clock is less than 15 minutes.

Since a wrong system time can cause serious problems (missed backups, dropped mail messages, mount failures on remote file systems, etc.), it is strongly recommended to *always* set the hardware clock to UTC.

POWER, AMD/Intel If a network is already configured, you can configure time synchronization with an NTP server. Click *Other Settings* to either alter the NTP settings or to *Manually* set the time. See Book “Administration Guide”, Chapter 38 “Time synchronization with NTP” for more information on configuring the NTP service. When finished, click *Accept* to continue the installation. ◀


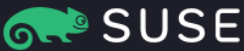
POWER, AMD/Intel If running without NTP configured, consider setting `SYST0HC=no` (`sysconfig` variable) to avoid saving unsynchronized time into the hardware clock. ◀

Note: Time cannot be changed on IBM Z

Since the operating system is not allowed to change time and date directly, the *Other Settings* option is not available on IBM Z.

9.13 Create new user

Create a local user in this step.



Local User

☒ Create New User

User's Full Name

Username

Password

Confirm Password

☐ Use this password for system administrator

☐ Automatic Login

☐ Skip User Creation

[Help](#) [Release Notes...](#) [Abort](#) [Back](#) [Next](#)

FIGURE 9.14: CREATE NEW USER

After entering the first name and last name, either accept the proposal or specify a new *User name* that will be used to log in. Only use lowercase letters (a-z), digits (0-9) and the characters `.` (dot), `-` (hyphen) and `_` (underscore). Special characters, umlauts and accented characters are not allowed.

Finally, enter a password for the user. Re-enter it for confirmation (to ensure that you did not type something else by mistake). To provide effective security, a password should be at least six characters long and consist of uppercase and lowercase letters, numbers and special characters (7-bit ASCII). Umlauts or accented characters are not allowed. Passwords you enter are checked for weakness. When entering a password that is easy to guess (such as a dictionary word or a name) you will see a warning. It is a good security practice to use strong passwords.



Important: User name and password

Remember both your user name and the password because they are needed each time you log in to the system.

If you install SUSE Linux Enterprise Server on a machine with one or more existing Linux installations, YaST allows you to import user data such as user names and passwords. Select *Import User Data from a Previous Installation* and then *Choose Users* for import.

If you do not want to configure any local users (for example when setting up a client on a network with centralized user authentication), skip this step by choosing *Next* and confirming the warning. Network user authentication can be configured at any time later in the installed system; refer to Book “Administration Guide”, Chapter 6 “Managing users with YaST” for instructions.

Two additional options are available:

Use this password for system administrator

If checked, the same password you have entered for the user will be used for the system administrator root. This option is suitable for stand-alone workstations or machines in a home network that are administrated by a single user. When not checked, you are prompted for a system administrator password in the next step of the installation workflow (see [Section 9.14, “Authentication for the system administrator root”](#)).

Automatic login

This option automatically logs the current user in to the system when it starts. This is mainly useful if the computer is operated by only one user. For automatic login to work, the option must be explicitly enabled.



Warning: Automatic login

With the automatic login enabled, the system boots straight into your desktop with no authentication. If you store sensitive data on your system, you should not enable this option if the computer can also be accessed by others.

In an environment where users are centrally managed (for example by NIS or LDAP) you should skip the creation of local users. Select *Skip User Creation* in this case.

9.14 Authentication for the system administrator root

If you have not chosen *Use this Password for System Administrator* in the previous step, you will be prompted to enter a password for the system administrator root or provide a public SSH key. Otherwise, this configuration step is skipped.


The screenshot shows the SUSE logo at the top left. The main heading is "Authentication for the System Administrator 'root'". Below this, there is a warning: "Do not forget what you enter here." followed by three input fields: "Password for root User", "Confirm Password", and "Test Keyboard Layout". Below these fields is the "Import Public SSH Key" section, which includes a dropdown menu showing "QEMU DVD-ROM (/dev/sr0)", a "Refresh" button, and a "Browse..." button. At the bottom, there are four buttons: "Help", "Release Notes...", "Abort", and "Next".

FIGURE 9.15: AUTHENTICATION FOR THE SYSTEM ADMINISTRATOR root

Enter the password for the system administrator root. For verification purposes, the password for root must be entered twice. Do not forget the password as it cannot be retrieved later.



Tip: Passwords and keyboard layout

It is recommended to only use US ASCII characters. In case of a system error or when you need to start your system in rescue mode, the keyboard may not be localized.

To change the root password later in the installed system, run YaST and start *Security and Users > User and Group Management*.



Important: The root user

root is the name of the system administrator or superuser. Its user ID (uid) is 0. Unlike regular users, the root account has unlimited privileges.

Do not forget the root password

Only root has the privileges to change the system configuration, install programs, manage users and set up new hardware. To carry out such tasks, the root password is required. Do not forget the password as it cannot be retrieved later.

Do not use the root user for daily work

Logging in as root for daily work is rather risky: Commands from root are usually executed without additional confirmation, so a single mistake can lead to an irretrievable loss of system files. Only use the root account for system administration, maintenance and repair.

Do not rename the root user account

YaST will always name the system administrator root. While it is technically possible to rename the root account, certain applications, scripts or third-party products may rely on the existence of a user called root. While such a configuration always targets individual environments, necessary adjustments could be overwritten by vendor updates, so this becomes an ongoing task rather than a one-time setting. This is especially true in very complex setups involving third-party applications, where it needs to be verified with every vendor involved whether a rename of the root account is supported.

As the implications for renaming the root account cannot be foreseen, SUSE does not support renaming the root account.

Usually, the idea behind renaming the root account is to hide it or make it unpredictable. However, `/etc/passwd` requires `644` permissions for regular users, so any user of the system can retrieve the login name for the user ID 0. For better ways to secure the root account, refer to *Book "Security and Hardening Guide", Chapter 14 "User management", Section 14.5 "Restricting root logins"* and *Book "Security and Hardening Guide", Chapter 14 "User management", Section 14.5.3 "Restricting SSH logins"*.

If you want to access the system remotely via SSH using a public key, import a key from a removable storage device or an existing partition. After the installation is finished, you can log in through SSH using the provided SSH key.

PROCEDURE 9.1: **ADDING A PUBLIC SSH KEY FOR USER `root`**

To import a public SSH key from a medium partition, perform the following steps:

1. The public SSH key is located in your `~/.ssh` directory and has the file extension `.pub`. Copy it to a removable storage device or an existing partition that is not formatted during installation.
2. If your key is on a removable storage device, insert it into your computer and click *Refresh*. You should see the device in the drop-down box under *Import Public Key*.
3. Click *Browse*, select the public SSH key and confirm with *Open*.
4. Proceed with *Next*.

If you have both set a password and added a public SSH key, and need remote access right after the installation, do not forget to open the SSH port in the *Security* section of the *Installation Settings* summary. If you set no password but only add a key, the port will be opened automatically to prevent you from being locked out of the newly installed system.

9.15 Installation settings

On the last step before the real installation takes place, you can alter installation settings suggested by the installer. To modify the suggestions, click the respective headline. After having made changes to a particular setting, you are always returned to the Installation Settings window, which is updated accordingly.

If you have added an SSH key for your `root` as mentioned in [Procedure 9.1](#), make sure to open the SSH port in the *Security* settings.

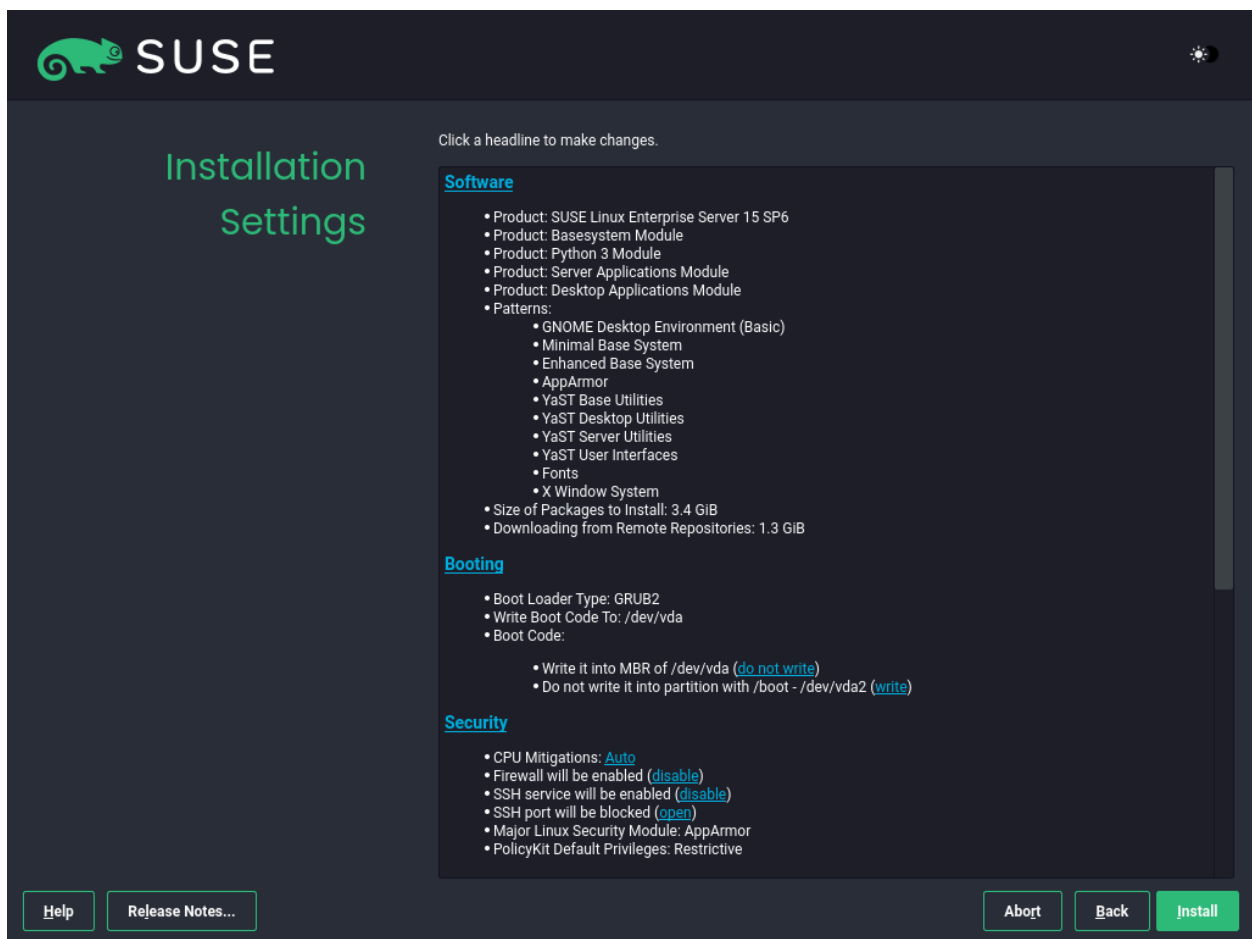


FIGURE 9.16: INSTALLATION SETTINGS

9.15.1 Software

SUSE Linux Enterprise Server contains several software patterns for various application purposes. The available choice of patterns and packages depends on your selection of modules and extensions.

Click *Software* to open the *Software Selection and System Tasks* screen where you can modify the pattern selection according to your needs. Select a pattern from the list and see a description in the right-hand part of the window.

Each pattern contains several software packages needed for specific functions (for example Web and LAMP server or a print server). For a more detailed selection based on software packages to install, select *Details* to switch to the YaST Software Manager.




FIGURE 9.17: SOFTWARE SELECTION AND SYSTEM TASKS

You can also install additional software packages or remove software packages from your system at any later time with the YaST Software Manager. For more information, refer to *Book "Administration Guide", Chapter 8 "Installing or removing software"*.

If you choose to install GNOME, SUSE Linux Enterprise Server is installed with the X.org display server. As an alternative to GNOME, the lightweight window manager IceWM can be installed. Select *Details* from the *Software Selection and System Tasks* screen and search for icewm.



Tip: IBM Z: Hardware cryptography support

The hardware cryptography stack is not installed by default. To install it, select *System z HW crypto support* in the *Software Selection and System Tasks* screen.



Tip: Adding secondary languages

The language you selected with the first step of the installation will be used as the primary (default) language for the system. You can add secondary languages from within the *Software* dialog by choosing *Details > View > Languages*.

9.15.2 Booting

The installer proposes a boot configuration for your system. Other operating systems found on your computer, such as Microsoft Windows or other Linux installations, will automatically be detected and added to the boot loader. However, SUSE Linux Enterprise Server will be booted by default. Normally, you can leave these settings unchanged. If you need a custom setup, modify the proposal according to your needs. For information, see *Book "Administration Guide", Chapter 18 "The boot loader GRUB 2", Section 18.3 "Configuring the boot loader with YaST"*.



Important: Software RAID 1

Booting a configuration where /boot resides on a software RAID 1 device is supported, but it requires to install the boot loader into the MBR (*Boot Loader Location > Boot from Master Boot Record*). Having /boot on software RAID devices with a level other than RAID 1 is not supported. Also see *Book "Storage Administration Guide", Chapter 8 "Configuring software RAID for the root partition"*.

9.15.3 Security

The *CPU Mitigations* refer to kernel boot command line parameters for software mitigations that have been deployed to prevent CPU side-channel attacks. Click the selected entry to choose a different option. For details, see *Book "Administration Guide", Chapter 18 "The boot loader GRUB 2" CPU Mitigations*.

By default, the *Firewall* is enabled on all configured network interfaces. To completely disable firewalld, click *disable* (not recommended).



Note: Firewall settings

When the firewall is activated, all interfaces are assigned to the public zone, where all ports are closed by default, ensuring maximum security. The only port you can open during the installation is port 22 (SSH), to allow remote access. Other services requiring network access (such as FTP, Samba, Web server, etc.) will only work after having adjusted the firewall settings. Refer to Book “*Security and Hardening Guide*”, Chapter 23 “*Masquerading and firewalls*” for configuration details.



Note: Firewall settings for receiving updates

By default, the firewall on SUSE Linux Enterprise Server only blocks incoming connections. If your system is behind another firewall that blocks outgoing traffic, make sure to allow connections to <https://scc.suse.com/> and <https://updates.suse.com> on ports 80 and 443 in order to receive updates.

The *SSH service* is enabled by default, but its port (22) is closed in the firewall. Click *open* to open the port or *disable* to disable the service. Note that if SSH is disabled, remote logins will not be possible. Refer to Book “*Security and Hardening Guide*”, Chapter 22 “*Securing network operations with OpenSSH*” for more information.



Tip: Existing SSH host keys

If you install SUSE Linux Enterprise Server on a machine with existing Linux installations, the installation routine imports an SSH host key. It chooses the host key with the most recent access time by default. See also [Section 9.15.9, “Import SSH host keys and configuration”](#).

If you are performing a remote administration over VNC, you can also specify whether the machine should be accessible via VNC after the installation. Note that enabling VNC also requires you to set the *Default systemd Target* to *graphical*.

The default *Major Linux Security Module* is *AppArmor*. To disable it, select *None* as module in the *Security* settings. This allows you to deselect the *AppArmor* pattern in the *Software* settings ([Section 9.15.1, “Software”](#)).

9.15.4 *Security Profiles*



Important: Availability in SUSE Linux Enterprise 15 SP4

This feature is available for SUSE Linux Enterprise 15 SP4 GM via installer self-update or using the QU2 media.

This category allows hardening your system with OpenSCAP security policies. The first policy that was implemented is the Security Technical Implementation Guide (STIG) by the Defense Information Systems Agency (DISA).

Click to *enable* the security policy. Non-compliant installation settings will be listed with the rule they violate. Some settings can be adjusted automatically by clicking *fix rule*. For settings that require user input, click *modify settings* to open the respective settings screen.



Tip: Checking policy compliance during installation

If you do not want to wait for the *Installation Settings* screen, but want the installer to check the settings from the beginning of the installation process, boot the system with the boot parameter YAST_SECURITY_POLICY=POLICY. To check for compliance with the DISA STIG, use YAST_SECURITY_POLICY=stig. For more information about boot parameters, refer to [Chapter 8, Boot parameters](#).

The installer does not check all rules of the profile, only those necessary for the installation or that are hard to fix afterward. To apply the remaining rules, a full SCAP remediation is performed on first boot. You can also perform a *scan only* or *do nothing* and manually remediate the system later with OpenSCAP. For more information, refer to the articles [Hardening SUSE Linux Enterprise with STIG](https://documentation.suse.com/compliance/all/html/SLES-stig/article-stig.html) (<https://documentation.suse.com/compliance/all/html/SLES-stig/article-stig.html>) [↗](#) and [Hardening SUSE Linux Enterprise with OpenSCAP](https://documentation.suse.com/compliance/all/html/SLES-openscap/article-openscap.html) (<https://documentation.suse.com/compliance/all/html/SLES-openscap/article-openscap.html>) [↗](#).

9.15.5 *Network configuration*

This category displays the current network settings, as automatically configured after booting into the installation (see [Section 9.6](#)) or as manually configured during the installation process. By default, wicked is used for server installations and NetworkManager for desktop workloads.

If you want to check or adjust the network settings, click *Network Configuration*. This takes you to the YaST *Network Settings* module. For details, see Book “Administration Guide”, Chapter 23 “Basic networking”, Section 23.4 “Configuring a network connection with YaST”.



Important: Support for NetworkManager

SUSE only supports NetworkManager for desktop workloads with SLED or the Workstation extension. All server certifications are done with **wicked** as the network configuration tool, and using NetworkManager may invalidate them. NetworkManager is not supported by SUSE for server workloads.

9.15.6 *Kdump*

Using Kdump, you can save a dump of the kernel (in case of a crash) to analyze what went wrong. Use this dialog to enable and configure Kdump. Find detailed information at Book “System Analysis and Tuning Guide”, Chapter 20 “Kexec and Kdump”.

9.15.7 *IBM Z: blacklist devices*

To save memory, all channels for devices currently not in use are blacklisted by default (each channel that is not blacklisted occupies approximately 50 KB of memory). To configure additional hardware in the installed system using channels that are currently blacklisted, run the respective YaST module to enable the respective channels first.

To disable blacklisting, click *disable*.

9.15.8 *Default systemd target*

SUSE Linux Enterprise Server can boot into two different targets (formerly known as “run-levels”). The *graphical* target starts a display manager, whereas the *multi-user* target starts the command line interface.

The default target is *graphical*. In case you have not installed the *X Window System* patterns, you need to change it to *multi-user*. If the system should be accessible via VNC, you need to choose *graphical*.

9.15.9 *Import SSH host keys and configuration*

If an existing Linux installation on your computer was detected, YaST will import the most recent SSH host key found in `/etc/ssh` by default, optionally including other files in the directory as well. This makes it possible to reuse the SSH identity of the existing installation, avoiding the `REMOTE HOST IDENTIFICATION HAS CHANGED` warning on the first connection. Note that this item is not shown in the installation summary if YaST has not discovered any other installations. You have the following choices:

I would like to import SSH keys from a previous install:

Select this option to import the SSH host key and optionally the configuration of an installed system. You can select the installation to import from in the option list below.

Import SSH Configuration

Enable this to copy other files in `/etc/ssh` to the installed system in addition to the host keys.

9.15.10 *System*

This screen lists all the hardware information the installer could obtain about your computer. When opened for the first time, the hardware detection is started. Depending on your system, this may take some time. Select any item in the list and click *Details* to see detailed information about the selected item. Use *Save to File* to save a detailed list to either the local file system or a removable device.

Advanced users can also change the *PCI ID Setup* and kernel settings by choosing *Kernel Settings*. A screen with two tabs opens:

PCI ID setup

Each kernel driver contains a list of device IDs of all devices it supports. If a new device is not in any driver's database, the device is treated as unsupported, even if it can be used with an existing driver. You can add PCI IDs to a device driver here. Only advanced users should attempt to do so.

To add an ID, click *Add* and select whether to *Manually* enter the data, or whether to choose from a list. Enter the required data. The *SysFS Dir* is the directory name from `/sys/bus/pci/drivers`—if empty, the *driver* name is used as the directory name. Existing entries can be managed with *Edit* and *Delete*.

Kernel settings

Change the *Global I/O Scheduler* here. If *Not Configured* is chosen, the default setting for the respective architecture will be used. This setting can also be changed at any time later from the installed system. Refer to Book “*System Analysis and Tuning Guide*”, Chapter 14 “*Tuning I/O performance*” for details on I/O tuning.

Also activate the *Enable SysRq Keys* here. These keys will let you issue basic commands (such as rebooting the system or writing kernel dumps) in case the system crashes. Enabling these keys is recommended when doing kernel development. Refer to <https://www.kernel.org/doc/html/latest/admin-guide/sysrq.html> for details.

9.16 Performing the installation

After configuring all installation settings, click *Install* in the Installation Settings window to start the installation. Some software may require a license confirmation. If your software selection includes such software, license confirmation dialogs are displayed. Click *Accept* to install the software package. When not agreeing to the license, click *I Disagree* and the software package will not be installed. In the dialog that follows, confirm with *Install* again.

The installation usually takes between 15 and 30 minutes, depending on the system performance and the selected software scope. After having prepared the hard disk and having saved and restored the user settings, the software installation starts. Choose *Details* to switch to the installation log or *Release Notes* to read important up-to-date information that was not available when the manuals were printed.

After the software installation has completed, the system reboots into the new installation where you can log in. To customize the system configuration or to install additional software packages, start YaST.

9.16.1 IBM Z: IPLing the installed system

YaST usually reboots into the installed system on the IBM Z platform. Exceptions are installations where the boot loader resides on an FCP device in environments with LPAR on a machine older than z196 or with z/VM older than release 5.4. The boot loader gets written to a separate partition mounted as /boot/zipl/.

In cases where an automatic reboot is not possible, YaST will show a dialog containing information about from which device to do an IPL. Accept the shutdown option and perform an IPL after the shutdown. The procedure varies according to the type of installation:

LPAR installation

In the IBM Z HMC, select *Load*, select *Clear*, then enter the loading address (the address of the device containing the `/boot/zipl` directory with the boot loader). If using a zFCP disk as the boot device, choose *Load from SCSI* and specify the load address of your FCP adapter plus WWPN and LUN of the boot device. Now start the loading process.

z/VM installation

Log in to the VM guest (see [Example 5.1, “Configuration of a z/VM directory”](#) for the configuration) as `LINUX1` and proceed to IPL the installed system:

```
IPL 151 CLEAR
```

`151` is an example address of the DASD boot device, replace this value with the correct address.

If using a zFCP disk as the boot device, specify both the zFCP WWPN and LUN of the boot device before initiating the IPL. The parameter length is limited to eight characters. Longer numbers must be separated by spaces:

```
SET LOADDEV PORT 50050763 00C590A9 LUN 50010000 00000000
```

Finally, initiate the IPL:

```
IPL FC00
```

`FC00` is an example address of the zFCP adapter, replace this value with the correct address.

KVM guest installation

After the installation has finished, the virtual machine is shut down. At this point, log in to the KVM host, edit the virtual machine's description file and restart it to IPL into the installed system:

1. Log in to the KVM host.
2. Edit the domain XML file by running

```
> sudo virsh edit s12-1
```

and remove the following lines:

```
<!-- Boot kernel - remove 3 lines after successful installation -->  
<kernel>/var/lib/libvirt/images/s12-kernel.boot</kernel>  
<initrd>/var/lib/libvirt/images/s12-initrd.boot</initrd>  
<cmdline>linuxrcstderr=/dev/console</cmdline>
```

3. Restart the VM Guest to IPL into the installed system:

```
> sudo virsh start s12-1 --console
```



Note: `cio_ignore` is disabled for KVM installations

The kernel parameter `cio_ignore` prevents the kernel from looking at all the available hardware devices. However, for KVM guests, the hypervisor already takes care to only provide access to the correct devices. Therefore `cio_ignore` is disabled by default when installing a KVM guest (for z/VM and LPAR installations it is activated by default).

9.16.2 IBM Z: Connecting to the installed system

After IPLing the system, establish a connection via VNC, SSH, or X to log in to the installed system. Using either VNC or SSH is recommended. To customize the system configuration or to install additional software packages, start YaST.

9.16.2.1 Using VNC to connect

A message in the 3270 terminal asks you to connect to the Linux system using a VNC client. However, this message is easily missed, because it is mixed with kernel messages and the terminal process might quit before you notice the message. If nothing happens for five minutes, try to initiate a connection to the Linux system using a VNC viewer.

If you connect using a JavaScript-capable browser, enter the complete URL, consisting of the IP address of the installed system along with the port number, in the following fashion:

```
http://IP_OF_INSTALLED_SYSTEM:5801/
```

9.16.2.2 Using SSH to connect

A message in the 3270 terminal asks you to connect to the Linux system with an SSH client. This message is easily missed, however, because it is mixed with kernel messages and the terminal process might quit before you become aware of the message.

When the message appears, use SSH to log in to the Linux system as root. If the connection is denied or times out, wait for the login timeout to expire, then try again (this time depends on server settings).

9.16.2.3 Using X to connect

When IPLing the installed system, make sure that the X server used for the first phase of the installation is up and still available before booting from the DASD. YaST opens on this X server to finish the installation. Complications may arise if the system is booted up but unable to connect to the X server in a timely fashion.


10 Registering SUSE Linux Enterprise and managing modules/extensions

To get technical support and product updates, you need to register and activate SUSE Linux Enterprise Server with the SUSE Customer Center. It is recommended to register during the installation, since this will enable you to install the system with the latest updates and patches available. However, if you are offline or want to skip the registration step, you can register at any time later from the installed system.

Modules and extensions add features to your system and allow you to customize the system according to your needs. These components also need to be registered and can be managed with YaST or command line tools. For more details also refer to the *Article “Modules and Extensions Quick Start”*.



Note: SUSE account

Registering with the SUSE Customer Center requires a SUSE account. In case you do not have a SUSE account yet, go to the SUSE Customer Center home page (<https://sc-c.suse.com/> ) to create one.



Tip: Deregistering a system

To completely deregister a system including all modules and extensions use the command line tool **SUSEConnect**. Deregistering a system removes its entry on the registration server and removes all repositories for modules, extensions, and the product itself.

```
> sudo SUSEConnect -d
```

10.1 Registering during the installation

The easiest and recommended way to register is during the installation. It not only allows you to install the latest patch level of SUSE Linux Enterprise Server, but also gives you access to all modules and extensions without having to provide additional installation media. This also applies to all modules or extensions you install. For details on the registration process, refer to [Section 9.7, “Registration”](#).

If the system was successfully registered during installation, YaST adds online repositories provided by SUSE Customer Center. This prevents problems if local installation sources are no longer available and ensures that you always get the latest updates from the online repositories.

10.2 Registering during automated deployment

If you deploy your instances automatically using AutoYaST, you can register the system during the installation by providing the respective information in the AutoYaST control file. Refer to *Book “AutoYaST Guide”, Chapter 4 “Configuration and installation options”, Section 4.3 “System registration and extension selection”* for details.

10.3 Registering from the installed system

If you skipped the registration during the installation or want to re-register your system, you can do it at any time using the YaST module *Product Registration* or the command-line tool **SUSEConnect**.

10.3.1 Registering with YaST

To register the system start *YaST > Software > Product Registration*. First register SUSE Linux Enterprise Server, then choose the modules and extensions you want to make available.




Important: Modules and extensions

If you installed the system from the `SLE-15-SP6-Full-ARCH-GM-media1.iso` media and skipped the registration, make sure to register all the modules and extensions you have chosen during the installation. You will only receive security updates and patches for modules and extensions that have been registered.

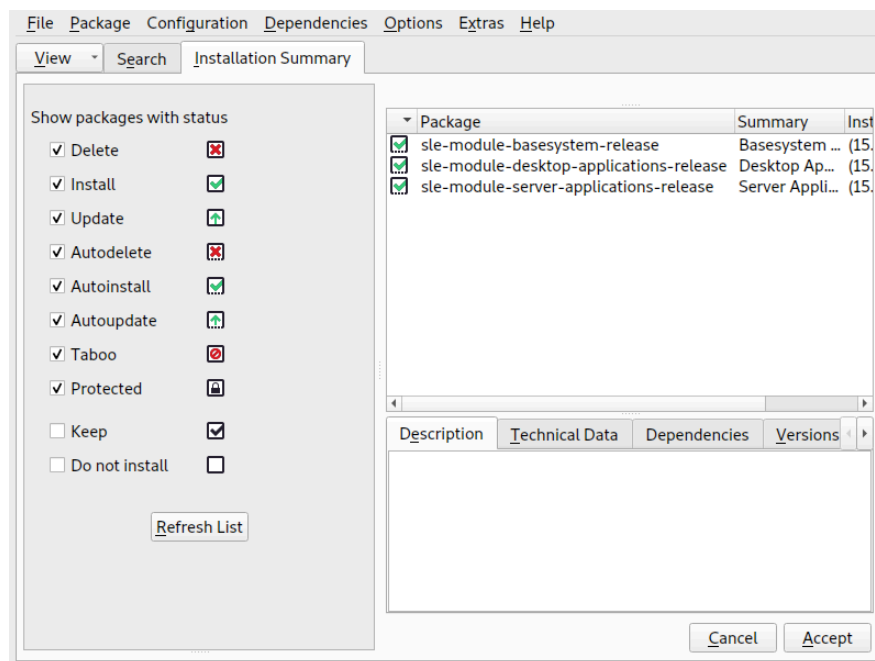
1. Start *YaST* > *Software* > *Product Registration*.

The screenshot shows a registration window titled "Registration" for "SUSE Linux Enterprise Server 15 SP3". It prompts the user to "Please select your preferred method of registration." There are two radio button options: "Register System via scc.suse.com" (which is selected) and "Register System via local RMT Server". Under the first option, there are input fields for "E-mail Address" and "Registration Code". Under the second option, there is a dropdown menu for "Local Registration Server URL" with the value "https://rmt.example.com". At the bottom, there are three buttons: "Help", "Abort", and "Next".

2. Provide the *E-mail address* associated with the SUSE account you or your organization uses to manage subscriptions. Also enter the *Registration Code* you received with your copy of SUSE Linux Enterprise Server.
3. By default, the system is registered with the SUSE Customer Center.
If your organization provides local registration servers, you can either choose one from the list of auto-detected servers or provide the URL at *Register System via local RMT Server*.
4. Choose *Next* to start the registration process. SUSE Linux Enterprise Server is registered with the chosen server and the associated repositories are added to your system. The *Extension and Module Selection* dialog opens.



5. Select all modules and extensions you want to make available in the system. At minimum, select the default modules (*Basesystem Module* and *Server Applications Module*). Also make sure to select any additional modules or extensions that you have added during the installation. Note that all extensions require additional registration codes that must be purchased. Proceed with *Next*.
6. Depending on your selection, you may need to accept one or more license agreements. All components registered with the chosen server and the associated repositories are added to your system.
7. The YaST package installer opens to install release-packages for each module and, depending on your choice of modules and extensions, additional packages. It is strongly recommended *not to deselect* any of the preselected packages; you may, however, add additional packages.



Choose *Accept* and *Finish* to conclude the registration process.

10.3.2 Registering with SUSEConnect

Registering the system, along with modules and extensions, can be done from the command line using **SUSEConnect**. For information on that topic, refer to the inline documentation with **man 8 SUSEConnect**

PROCEDURE 10.2: PRODUCT REGISTRATION WITH SUSECONNECT

1. To register SUSE Linux Enterprise Server with SUSE Customer Center run **SUSEConnect** as follows:

```
> sudo SUSEConnect -r REGISTRATION_CODE -e EMAIL_ADDRESS
```

To register with a local registration server, provide the URL of the server:

```
> sudo SUSEConnect -r REGISTRATION_CODE -e EMAIL_ADDRESS \
--url "https://suse_register.example.com/"
```

Replace REGISTRATION_CODE with the registration code you received with your copy of SUSE Linux Enterprise Server. Replace EMAIL_ADDRESS with the E-mail address associated with the SUSE account you or your organization uses to manage subscriptions.

This process will register the *Basesystem Module* and *Server Applications Module* and add the associated repositories to your system.

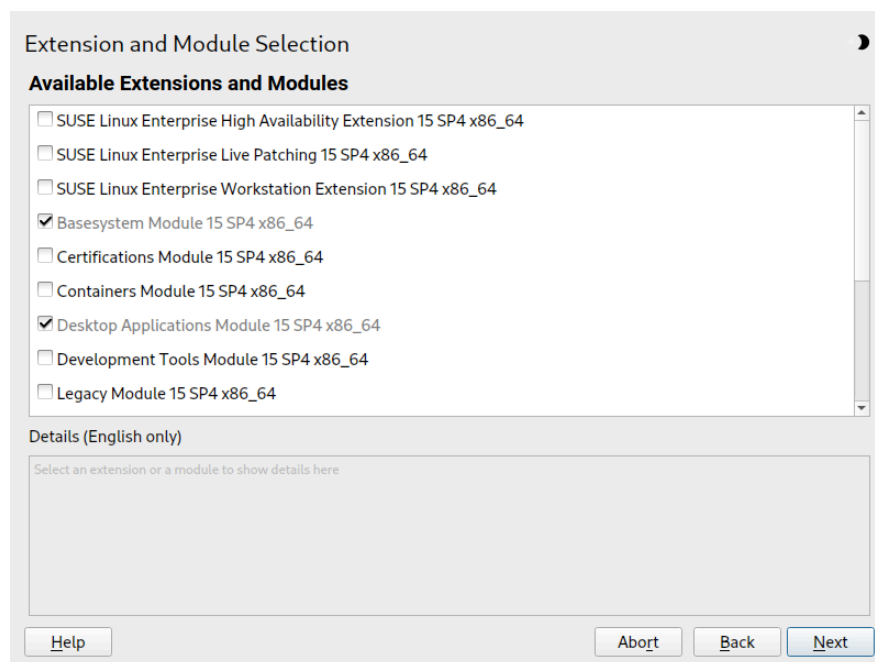
2. SUSE Linux Enterprise Server including the two default repositories is now registered. In case you want to register additional modules or extensions, proceed as outlined in [Section 10.4, “Managing modules and extensions in a running system”](#).

10.4 Managing modules and extensions in a running system

You can add and remove modules and extensions even after a system is installed and registered. You can use either YaST or **SUSEConnect** to do that. For additional information, refer to the *Article “Modules and Extensions Quick Start”*.

10.4.1 Adding modules and extensions with YaST

1. Start *YaST > Software > System Extensions*.



2. To add modules or extensions, select all components you want to install. Note that all extensions require additional registration codes.

3. All additional components are registered with the registration server and the associated repositories are added to your system.
4. The YaST package installer opens to install release-packages for each module and, depending on your choice of modules and extensions, additional packages. It is strongly recommended *not to deselect* any of the preselected packages; you may, however, add additional packages.

Choose *Accept* and *Finish* to conclude the process.




Tip: Module dependencies

Similar to software packages, which may depend on other packages to function, a module may have dependencies on other modules. If this is the case, the modules on which it depends are automatically selected for installation.

10.4.2 Deleting modules and extensions with YaST

1. Start *YaST > Software > Add-On Products*.



2. Choose the module or extension that should be removed and click *Delete*. Confirm the warning saying that all packages from the selected component will be removed.

3. The YaST Software Manager opens and lists all installed packages from the deleted module or extension. Click *Accept* to remove all of them. It is strongly recommended to do so, because you will no longer get updates for packages from deleted modules or extensions. In case you keep packages, make sure to at least remove the *-release package for each module or extension that gets deleted.

Proceed with *Accept* and then *OK*.



Warning: Deleting modules

Note that you should never delete the *Basesystem Module*. It is also not recommended to delete the *Server Applications Module*.



Warning: No updates for packages from deleted modules and extensions

If you choose to keep packages from deleted modules or extensions, you will no longer receive updates for these packages. Because this includes security fixes, keeping such packages may introduce a security risk to your system.

10.4.3 Adding or deleting modules and extensions with SUSEConnect

1. Run **SUSEConnect -list-extensions** to get an overview of available extensions:

```
> sudo SUSEConnect -list-extensions
AVAILABLE EXTENSIONS AND MODULES

Basesystem Module 15 SP6 x86_64 (Installed)
Deactivate with: SUSEConnect -d -p sle-module-basesystem/15.6/x86_64

Containers Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-containers/15.6/x86_64

Desktop Applications Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-desktop-applications/15.6/x86_64

Development Tools Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-development-tools/15.6/x86_64
```

```
SUSE Linux Enterprise Workstation Extension 15 SP6 x86_64
Activate with: SUSEConnect -p sle-we/15.6/x86_64 -r ADDITIONAL REGCODE
```

```
SUSE Cloud Application Platform Tools Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-cap-tools/15.6/x86_64
```

```
SUSE Linux Enterprise Live Patching 15 SP6 x86_64
Activate with:
    SUSEConnect -p sle-module-live-patching/15.6/x86_64 -r ADDITIONAL REGCODE
```

```
SUSE Package Hub 15 SP6 x86_64
Activate with: SUSEConnect -p PackageHub/15.6/x86_64
```

```
Server Applications Module 15 SP6 x86_64 (Installed)
Deactivate with: SUSEConnect -d -p sle-module-server-applications/15.6/x86_64
```

```
Legacy Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-legacy/15.6/x86_64
```

```
Public Cloud Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-public-cloud/15.6/x86_64
```

```
SUSE Enterprise Storage 6 x86_64
Activate with: SUSEConnect -p ses/6/x86_64 -r ADDITIONAL REGCODE
```

```
SUSE Linux Enterprise High Availability Extension 15 SP6 x86_64
Activate with: SUSEConnect -p sle-ha/15.6/x86_64 -r ADDITIONAL REGCODE
```

```
Web and Scripting Module 15 SP6 x86_64
Activate with: SUSEConnect -p sle-module-web-scripting/15.6/x86_64
```

MORE INFORMATION

You can find more information about available modules here:
<https://www.suse.com/products/server/features/modules.html>

2. Run the appropriate command to add or delete a component. Note that adding extensions requires additional registration codes.



Warning: Deleting modules

Do not delete the *Basesystem Module*. It is also not recommended to delete the *Server Applications Module*.

Important: No automatic installation or removal of packages

SUSEConnect only adds or removes modules and extensions. It registers or deregisters the components and enables or disables their repositories, but it does not install or remove any packages. If you want this to be done automatically, use YaST to manage modules and extensions.

When adding a module or extension, **SUSEConnect** does not install default packages or patterns. To do this manually, use Zypper or YaST › *Software Management*.

When deleting a module or extension, **SUSEConnect** does not perform a cleanup. Packages from the module or extension remain installed on the system, but are no longer updated from a repository. To list these “orphaned” packages, run **zypper packages --orphaned**. To remove one or more packages, run **zypper remove PACKAGE [ANOTHER_PACKAGE]**. Alternatively use YaST › *Software Management* and then *View › Package Classification › Orphaned Packages* to list and delete orphaned packages.

Warning: No updates for packages from deleted modules and extensions

If you choose to keep packages from deleted modules or extensions, you will no longer receive updates for these packages. Because this includes security fixes, keeping such packages may introduce a security risk to your system.

10.5 SUSEConnect keep-alive timer

From version 0.3.33, the SUSEConnect package ships with two systemd units:

- suseconnect-keepalive.service: a service which runs the command **SUSEConnect --keep-alive** on demand.
- suseconnect-keepalive.timer: a timer which runs the service suseconnect-keepalive.service once a day.

These units are responsible for keeping the system information up-to-date with the SUSE Customer Center or registration server, and to provide accurate data about subscription usage.

The command **SUSEConnect --keep-alive** updates the last time a system has been seen and its hardware information with the registration service.



Note: The timer is enabled automatically

When the SUSEConnect package is installed or updated, and its version is equal to or greater than the one described above, the keep-alive timer will be enabled automatically.



Tip: Disabling the SUSEConnect keep-alive timer

If you prefer to not have the SUSEConnect keep-alive timer running on your system, you can disable it with **systemctl**:

```
> sudo systemctl disable --now suseconnect-keepalive.timer
```

Once the timer is disabled, subsequent updates to the SUSEConnect package will not reenable it.

11 *Expert Partitioner*

Sophisticated system configurations require specific disk setups. You can perform all common partitioning tasks during the installation.

To get persistent device naming with block devices, use the block devices below `/dev/disk/by-id` or `/dev/disk/by-uuid`.

Logical Volume Management (LVM) is a disk partitioning scheme that is designed to be much more flexible than the physical partitioning used in standard setups. Its snapshot functionality enables easy creation of data backups. Redundant Array of Independent Disks (RAID) offers increased data integrity, performance, and fault tolerance. SUSE Linux Enterprise Server also supports multipath I/O (see Book “*Storage Administration Guide*”, Chapter 18 “*Managing multipath I/O for devices*” for details). There is also the option to use iSCSI as a networked disk (read more about iSCSI in Book “*Storage Administration Guide*”, Chapter 15 “*Mass storage over IP networks: iSCSI*”).



Warning: Disk space units

Note that for partitioning purposes, disk space is measured in binary units, rather than in decimal units. For example, if you enter sizes of 1GB, 1GiB or 1G, they all signify 1 GiB (Gibibyte), as opposed to 1 GB (Gigabyte).

Binary

1 GiB = 1 073 741 824 bytes.

Decimal

1 GB = 1 000 000 000 bytes.

Difference

1 GiB ≈ 1.07 GB.

11.1 *Using the Expert Partitioner*

Using the *Expert Partitioner* (Figure 11.1, “*The YaST partitioner*”), you can add, delete, resize, and edit partitions, as well as access the soft RAID, and LVM configuration.



Warning: Repartitioning the running system

Although it is possible to repartition your system while it is running, the risk of making a mistake that causes data loss is very high. Try to avoid repartitioning your installed system and always create a complete backup of your data before attempting to do so.




FIGURE 11.1: THE YAST PARTITIONER



Tip: IBM Z: Device names

IBM Z recognizes only DASD, zFCP and SCSI hard disks. IDE hard disks are not supported. This is why these devices appear in the partition table as dasda or sda for the first recognized device.

All existing or suggested partitions on all connected hard disks are displayed in the list of *Available Storage* in the YaST *Expert Partitioner* dialog. Entire hard disks are listed as devices without numbers, such as /dev/sda (or /dev/dasda). Partitions are listed as parts of these devices, such as /dev/sda1 (or /dev/dasda1, respectively). The size, type, encryption status, file system, and mount point of the hard disks and their partitions are also displayed. The mount point describes where the partition appears in the Linux file system tree.

Several functional views are available on the left hand *System View*. These views can be used to collect information about existing storage configurations, configure functions (like RAID, Volume Management, Crypt Files), and view file systems with additional features, such as Btrfs, NFS, or TMPFS.

If you run the expert dialog during installation, any free hard disk space is also listed and automatically selected. To provide more disk space to SUSE Linux Enterprise Server, free the needed space by going from the bottom toward the top in the list of partitions.

11.1.1 Partition tables

SUSE Linux Enterprise Server allows to use and create different *partition tables*. In some cases the partition table is called *disk label*. The partition table is important to the boot process of your computer. To boot your machine from a partition in a newly created partition table, make sure that the table format is supported by the firmware.

To change the partition table, click the relevant disk name in the *System View* and choose *Expert > Create New Partition Table*.

11.1.1.1 Master boot record

The *master boot record (MBR)* is the legacy partition table used on IBM PCs. It is sometimes also called an *MS-DOS* partition table. The MBR only supports four primary partitions. If the disk already has an MBR, SUSE Linux Enterprise Server allows you to create additional partitions in it which can be used as the installation target.

The limit of four partitions can be overcome by creating an *extended partition*. The extended partition itself is a primary partition and can contain more *logical partitions*.

UEFI firmware usually supports booting from MBR in the legacy mode.

11.1.1.2 GPT partition table

UEFI computers use a *GUID Partition Table (GPT)* by default. SUSE Linux Enterprise Server will create a GPT on a disk if no other partition table exists.

Old BIOS firmware does not support booting from GPT partitions.

You need a GPT partition table to use one of the following features:

- More than four primary partitions
- UEFI Secure Boot
- Use disks larger than 2 TB



Note: Mislabeled partitions created with Parted 3.1 or earlier versions

GPT partitions created with Parted 3.1 or earlier versions use the Microsoft Basic Data partition type instead of the newer Linux-specific GPT GUID. Newer versions of Parted set the misleading flag `msftdata` on such partitions. This causes various disk tools to label the partition as a *Windows Data Partition* or similar.

To remove the flag, run:

```
# parted DEVICE set PARTITION_NUMBER msftdata off
```

11.1.1.3 Partition tables on IBM Z

On IBM Z platforms, SUSE Linux Enterprise Server supports *SCSI hard disks* and *direct access storage devices* (DASD). While SCSI disks can be partitioned as described above, DASDs can have no more than three partition entries in their partition tables.

11.1.2 Partitions

The YaST Partitioner can create and format partitions with several file systems. The default file system used by SUSE Linux Enterprise Server is Btrfs. For details, see [Section 11.1.2.2, “Btrfs partitioning”](#).

Other commonly used file systems are available: Ext2, Ext3, Ext4, FAT, XFS, Swap, and UDF.

11.1.2.1 Creating a partition

To create a partition select *Hard Disks* and then a hard disk with free space. The actual modification can be done in the *Partitions* tab:

1. Click *Add* to create a new partition. When using *MBR*, specify to create a primary or extended partition. Within the extended partition, you can create several logical partitions. For details, see [Section 11.1.1, “Partition tables”](#).
2. Specify the size of the new partition. You can either choose to occupy all the free unpartitioned space, or enter a custom size.
3. Select the file system to use and a mount point. YaST suggests a mount point for each partition created. To use a different mount method, like mount by label, select *Fstab Options*.
4. Specify additional file system options if your setup requires them. This is necessary, for example, if you need persistent device names. For details on the available options, refer to [Section 11.1.3, “Editing a partition”](#).
5. Click *Finish* to apply your partitioning setup and leave the partitioning module.
If you created the partition during installation, you are returned to the installation overview screen.

11.1.2.2 Btrfs partitioning

The default file system for the root partition is Btrfs. For details, see *Book “Administration Guide”, Chapter 10 “System recovery and snapshot management with Snapper”* and *Book “Storage Administration Guide”, Chapter 1 “Overview of file systems in Linux”*. The root file system is the default subvolume and it is not listed in the list of created subvolumes. As a default Btrfs subvolume, it can be mounted as a normal file system.



Important: Btrfs on an encrypted root partition

The default partitioning setup suggests the root partition as Btrfs with `/boot` being a directory. To encrypt the root partition, make sure to use the GPT partition table type instead of the default MSDOS type. Otherwise the GRUB2 boot loader may not have enough space for the second stage loader.

It is possible to create snapshots of Btrfs subvolumes—either manually, or automatically based on system events. For example when making changes to the file system, **zypper** invokes the **snapper** command to create snapshots before and after the change. This is useful if you are not satisfied with the change **zypper** made and want to restore the previous state. As **snapper** invoked by **zypper** creates snapshots of the *root* file system by default, it makes sense to exclude specific directories from snapshots. This is the reason YaST suggests creating the following separate subvolumes:

/boot/grub2/i386-pc, /boot/grub2/x86_64-efi, /boot/grub2/powerpc-ieee1275, /boot/grub2/s390x-emu

A rollback of the boot loader configuration is not supported. The directories listed above are architecture-specific. The first two directories are present on AMD64/Intel 64 machines, the latter two on IBM POWER and on IBM Z, respectively.

/home

If /home does not reside on a separate partition, it is excluded to avoid data loss on rollbacks.

/opt

Third-party products usually get installed to /opt. It is excluded to avoid uninstalling these applications on rollbacks.

/srv

Contains data for Web and FTP servers. It is excluded to avoid data loss on rollbacks.

/tmp

All directories containing temporary files and caches are excluded from snapshots.

/usr/local

This directory is used when manually installing software. It is excluded to avoid uninstalling these installations on rollbacks.

/var

This directory contains many variable files, including logs, temporary caches, third party products in /var/opt, and is the default location for virtual machine images and databases. Therefore this subvolume is created to exclude all of this variable data from snapshots and has Copy-On-Write disabled.



Tip: Size of Btrfs partition

Since saved snapshots require more disk space, it is recommended to reserve enough space for Btrfs. While the minimum size for a root Btrfs partition with snapshots and default subvolumes is 16 GB, SUSE recommends at least 32 GB, or more if /home does not reside on a separate partition.

11.1.2.3 Managing Btrfs subvolumes using YaST

Subvolumes of a Btrfs partition can be now managed with the YaST *Expert Partitioner* module. You can add new or delete existing subvolumes.

PROCEDURE 11.1: BTRFS SUBVOLUMES WITH YAST

1. Choose *Btrfs* in the left side pane.
2. Select the Btrfs partition whose subvolumes you need to manage.
3. Depending on whether you want to edit, add, or delete subvolumes, do the following:
 - a. To edit a subvolume, select it from the list and click *Edit*. You can then disable copy-on-write (check *noCoW*) for the volume or limit its size. Click *Accept* to finish.
 - b. To add a new subvolume, click *Add Subvolume*, and enter its path. Optionally, you can disable copy-on-write (check *noCoW*) for the volume or limit its size. Click *Accept* to finish.
 - c. To delete a subvolume, select it from the list and click *Delete*. Confirm the deletion by clicking *Yes*.
 - d.

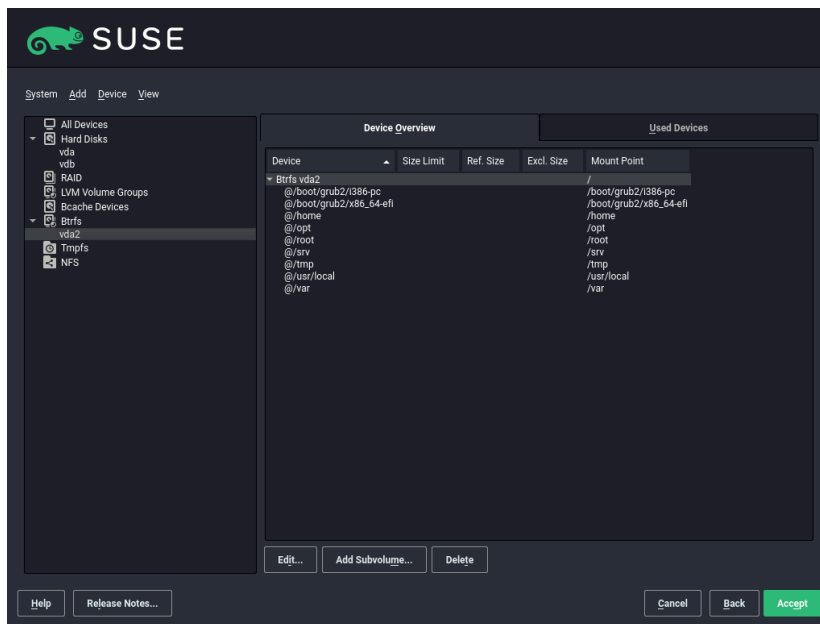


FIGURE 11.2: BTRFS SUBVOLUMES IN YAST PARTITIONER

4. Leave the partitioner with *Finish*.

11.1.3 Editing a partition

When you create a new partition or modify an existing partition, you can set various parameters. For new partitions, the default parameters set by YaST are usually sufficient and do not require any modification. To edit your partition setup manually, proceed as follows:

1. Select the partition.
2. Click *Edit* to edit the partition and set the parameters:

File system ID

Even if you do not want to format the partition at this stage, assign it a file system ID to ensure that the partition is registered correctly. Typical values are *Linux*, *Linux swap*, *Linux LVM*, and *Linux RAID*.

File System

To change the partition file system, click *Format Partition* and select file system type in the *File System* list.

SUSE Linux Enterprise Server supports several types of file systems. Btrfs is the Linux file system of choice for the root partition because of its advanced features. It supports copy-on-write functionality, creating snapshots, multi-device spanning, subvolumes, and other useful techniques. XFS, Ext3, and Ext4 are journaling file systems. These file systems can restore the system very quickly after a system crash, using write processes logged during the operation. Ext2 is not a journaling file system, but it is adequate for smaller partitions because it does not require much disk space for management.

The default file system for the root partition is Btrfs. The default file system for additional partitions is XFS.

The UDF file system can be used on optical rewritable and non-rewritable media, USB flash drives and hard disks. It is supported by multiple operating systems.

Swap is a special format that allows the partition to be used as a virtual memory. Create a swap partition of at least 256 MB. However, if you use up your swap space, consider adding memory to your system instead of adding swap space.



Warning: Changing the file system

Changing the file system and reformatting partitions irreversibly deletes all data from the partition.

For details on the various file systems, refer to Storage Administration Guide.

Encrypt Device

If you activate the encryption, all data is written to the hard disk in encrypted form. This increases the security of sensitive data, but reduces the system speed, as the encryption takes some time to process. More information about the encryption of file systems is provided in [Section 11.2, “Device encryption”](#) and Book “*Security and Hardening Guide*”, Chapter 12 “*Encrypting partitions and files*”.

Mount Point

Specify the directory where the partition should be mounted in the file system tree. Select from YaST suggestions or enter any other name.

Fstab Options

Specify various parameters contained in the global file system administration file (/etc/fstab). The default settings should suffice for most setups. You can, for example, change the file system identification from the device name to a volume label. In the volume label, use all characters except / and space.

To get persistent devices names, use the mount option *Device ID*, *UUID* or *LABEL*. In SUSE Linux Enterprise Server, persistent device names are enabled by default.



Note: IBM Z: Mounting by path

Since mounting by ID causes problems on IBM Z when using disk-to-disk copying for cloning purposes, devices are mounted by path in /etc/fstab on IBM Z by default.

If you prefer to mount the partition by its label, you need to define one in the *Volume label* text entry. For example, you could use the partition label HOME for a partition intended to mount to /home.

If you intend to use quotas on the file system, use the mount option *Enable Quota Support*. This must be done before you can define quotas for users in the YaST *User Management* module. For further information on how to configure user quota, refer to Book “Administration Guide”, Chapter 6 “Managing users with YaST”, Section 6.3.3 “Managing quotas”.

If you intend to specify quotas for Btrfs subvolumes, refer to Book “Storage Administration Guide”, Chapter 1 “Overview of file systems in Linux”, Section 1.2.5 “Btrfs quota support for subvolumes”.

3. Select *Finish* to save the changes.



Note: Resize file systems

To resize an existing file system, select the partition and use *Resize*. Note, that it is not possible to resize partitions while mounted. To resize partitions, unmount the relevant partition before running the partitioner.

11.1.4 Expert options

After you select a hard disk device (like *sda*) in the *System View* pane, you can access the *Expert* menu in the lower right part of the *Expert Partitioner* window. The menu contains the following commands:

Create new partition table

This option helps you create a new partition table on the selected device.



Warning: Creating a new partition table

Creating a new partition table on a device irreversibly deletes all partitions and their data from that device.

Clone this disk

This option helps you clone the device partition layout (but not the data) to other available disk devices.

11.1.5 Advanced options

After you select the host name of the computer (the top-level of the tree in the *System View* pane), you can access the *Configure* menu in the lower right part of the *Expert Partitioner* window. The menu contains the following commands:

Configure iSCSI

To access SCSI over IP block devices, you first need to configure iSCSI. This results in additionally available devices in the main partition list.

Configure multipath

Selecting this option helps you configure the multipath enhancement to the supported mass storage devices.

11.1.6 More partitioning tips

The following section includes a few hints and tips on partitioning that should help you make the right decisions when setting up your system.

11.1.6.1 Cylinder numbers

Note, that different partitioning tools may start counting the cylinders of a partition with 0 or with 1. When calculating the number of cylinders, you should always use the difference between the last and the first cylinder number and add one.

11.1.6.2 Using swap

Swap is used to extend the available physical memory. It is then possible to use more memory than physical RAM available. The memory management system of kernels before 2.4.10 needed swap as a safety measure. Then, if you did not have twice the size of your RAM in swap, the performance of the system suffered. These limitations no longer exist.

Linux uses a page called “Least Recently Used” (LRU) to select pages that might be moved from memory to disk. Therefore, running applications have more memory available and caching works more smoothly.

If an application tries to allocate the maximum allowed memory, problems with swap can arise. There are three major scenarios to look at:

System with no swap

The application gets the maximum allowed memory. All caches are freed, and thus all other running applications are slowed. After a few minutes, the kernel's out-of-memory kill mechanism activates and kills the process.

System with medium sized swap (128 MB–512 MB)

At first, the system slows like a system without swap. After all physical RAM has been allocated, swap space is used as well. At this point, the system becomes very slow and it becomes impossible to run commands from remote. Depending on the speed of the hard disks that run the swap space, the system stays in this condition for about 10 to 15 minutes until the out-of-memory kill mechanism resolves the issue. Note that you will need a certain amount of swap if the computer needs to perform a “suspend to disk”. In that case, the swap size should be large enough to contain the necessary data from memory (512 MB–1GB).

System with lots of swap (several GB)

It is better to not have an application that is out of control and swapping excessively in this case. If you use such application, the system will need many hours to recover. In the process, it is likely that other processes get timeouts and faults, leaving the system in an undefined state, even after terminating the faulty process. In this case, do a hard machine reboot and try to get it running again. Lots of swap is only useful if you have

an application that relies on this feature. Such applications (like databases or graphics manipulation programs) often have an option to directly use hard disk space for their needs. It is advisable to use this option instead of using lots of swap space.

If your system is not out of control, but needs more swap after some time, it is possible to extend the swap space online. If you prepared a partition for swap space, add this partition with YaST. If you do not have a partition available, you can also use a swap file to extend the swap. Swap files are generally slower than partitions, but compared to physical RAM, both are extremely slow so the actual difference is negligible.

PROCEDURE 11.2: ADDING A SWAP FILE MANUALLY

To add a swap file in the running system, proceed as follows:

1. Create an empty file in your system. For example, to add a swap file with 128 MB swap at `/var/lib/swap/swapfile`, use the commands:

```
> sudo mkdir -p /var/lib/swap
> sudo dd if=/dev/zero of=/var/lib/swap/swapfile bs=1M count=128
```

2. Initialize this swap file with the command

```
> sudo mkswap /var/lib/swap/swapfile
```



Note: Changed UUID for swap partitions when formatting via **mkswap**

Do not reformat existing swap partitions with **mkswap** if possible. Reformatting with **mkswap** will change the UUID value of the swap partition. Either reformat via YaST (which will update `/etc/fstab`) or adjust `/etc/fstab` manually.

3. Activate the swap with the command

```
> sudo swapon /var/lib/swap/swapfile
```

To disable this swap file, use the command

```
> sudo swapoff /var/lib/swap/swapfile
```

4. Check the current available swap spaces with the command

```
> cat /proc/swaps
```

Note that at this point, it is only temporary swap space. After the next reboot, it is no longer used.

5. To enable this swap file permanently, add the following line to `/etc/fstab`:

```
/var/lib/swap/swapfile swap swap defaults 0 0
```

11.1.7 Partitioning and LVM

From the *Expert Partitioner*, access the LVM configuration by clicking the *Volume Management* item in the *System View* pane. However, if a working LVM configuration already exists on your system, it is automatically activated upon entering the initial LVM configuration of a session. In this case, all disks containing a partition (belonging to an activated volume group) cannot be repartitioned. The Linux kernel cannot reread the modified partition table of a hard disk when any partition on this disk is in use. If you already have a working LVM configuration on your system, physical repartitioning should not be necessary. Instead, change the configuration of the logical volumes.

At the beginning of the physical volumes (PVs), information about the volume is written to the partition. To reuse such a partition for other non-LVM purposes, it is advisable to delete the beginning of this volume. For example, in the VG `system` and PV `/dev/sda2`, do this with the command:

```
dd if=/dev/zero of=/dev/sda2 bs=512 count=1
```



Warning: File system for booting

The file system used for booting (the root file system or `/boot`) must not be stored on an LVM logical volume. Instead, store it on a normal physical partition.

For more details about LVM, see Book “*Storage Administration Guide*”.

11.2 Device encryption

Linux Unified Key Setup (LUKS) is the standard for Linux disk encryption. It provides a standardized on-disk format and enables users to transport or migrate data seamlessly.

LUKS is used to encrypt block devices. The contents of the encrypted device are arbitrary, and therefore any file system can be encrypted, including swap partitions. All necessary setup information, like encryption keys and parameters, such as cipher type and key size, is stored in the partition header.

Encryption is done with a multi-layer approach. First, the block device is encrypted using a master key. Then, this master key is encrypted with each active user keys. User keys are derived from passphrases, FIDO2 security keys, TPMs or smart cards. This multi-layer approach allows users to change their passphrase without re-encrypting the whole block device.

For more information about LUKS, refer to *Book "Security and Hardening Guide", Chapter 13 "Storage encryption for hosted applications with cryptctl"*.

11.2.1 Encryption methods

To encrypt a device, follow the instructions in [Section 11.1.3, "Editing a partition"](#).



Tip: Enabling LUKS2 support in YaST

LUKS2 encryption is supported by the YaST Partitioner as of SUSE Linux Enterprise 15 SP4, but needs to be enabled explicitly. There are two ways to do this:

1. At boot time, by adding the parameter to `YAST_LUKS2_AVAILABLE` to the kernel command line. For information about boot parameters, refer to [Chapter 8, Boot parameters](#).
2. During installation in the YaST configuration:
 - In the graphical interface, press `Ctrl – Alt – Shift – C`.
 - In the text interface, press `Ctrl – D` and then `Shift – C`.

Check *Enable Experimental LUKS2 Encryption Support* and exit the configuration screen with `OK`.

If you do not enable LUKS2 support, the *Encryption method* selection is not visible and you only need to enter the encryption password.

Regular LUKS1

This method allows to encrypt the device using LUKS1. You have to provide the encryption password. Additional passwords—up to eight in total—can be added later with `cryptsetup luksAddKey`.

Regular LUKS2

LUKS2 uses a newer version of the header format, which is resilient to corruption, and supports up to 32 user keys and device labels. You have to provide the encryption password and the password-based key derivation function (PBKDF) that will be used to protect that passphrase (see [Section 11.2.2, “Password-based key derivation functions”](#)).

Pervasive LUKS2 (only on IBM Z)

This method allows to encrypt the device using LUKS2 with a master secure key processed by a Crypto Express cryptographic coprocessor configured in CCA mode. If the cryptographic system already contains a secure key associated to this volume, that key will be used. Otherwise, a new secure key will be generated and registered in the system. You need to provide an encryption password that will be used to protect the access to that master key. Moreover, when there are several APQNs in the system, you can select which ones to use.

For more information about pervasive encryption, refer to <https://www.ibm.com/docs/en/linux-on-systems?topic=security-pervasive-encryption>.

Encryption with Volatile Random Key (only for swap devices)

This method encrypts a swap device with a randomly generated key at boot and therefore does not support hibernation to hard disk. The swap device is re-encrypted on every boot, and its previous content is destroyed. To avoid data loss, disable hibernation and configure your system to shut down instead.

In addition to the encryption key, the device label and the UUID change every time the swap is re-encrypted, so neither is a valid option to mount a randomly encrypted swap device. Make sure the swap device is referenced by a stable name that is not subject to change on every reboot in the `/etc/crypttab` file. For example, for a swap partition it is safer to use the udev device id or path instead of the partition device name, since that device name may be assigned to a different partition during the next boot. If that happens, a wrong device could be encrypted instead of your swap!

YaST tries to use stable names in `/etc/crypttab`, unless it is configured to always use device names (see the *Settings* section of the partitioner). But for some devices, finding a fully stable name may not be possible. Only use encryption with volatile keys if you are sure about the implications.

Protected Swap (only for swap devices)

This method encrypts a swap device with a volatile protected AES key without requiring a cryptographic coprocessor. It is an improvement over the Encryption with Volatile Random Key method and all considerations for that method still apply.


Secure Swap (only for swap devices)

This method encrypts a swap device with a volatile secure AES key generated from a cryptographic coprocessor. It is an improvement over the Encryption with Volatile Random Key method and all considerations for that method still apply.


11.2.2 Password-based key derivation functions

The password-based key derivation function (PBKDF) to use depends on the context, the hardware capabilities and the needed level of compatibility with other system components:

PBKDF2

PBKDF2 is the function that LUKS1 uses. It is defined in RFC 2898 (<https://tools.ietf.org/html/rfc2898>) .

Argon2i

Argon2 is a function designed to be more secure and to require a lot of memory to be computed. It is defined in RFC 9106 (<https://tools.ietf.org/html/rfc9106>) . Argon2i is a variant of Argon2 optimized to resist side-channel attacks by accessing the memory array in a password-independent order.

Argon2id

Argon2id is a hybrid version of Argon2. It follows the Argon2i approach for the first half pass over memory and the Argon2d (not supported by YaST) approach to limit GPU cracking attacks for subsequent passes. RFC 9106 recommends using Argon2id if you do not know the difference between the types or you consider side-channel attacks to be a viable threat.

While [Argon2](#) is more secure, there are still use cases for [PBKDF2](#):

- As an intentional security feature, Argon2 requires a lot more memory to be computed. This may result in problems on some systems. If the strength of the password can be fully assured, then using PBKDF2 may still be secure and save memory.
- [grub2](#) offers limited support to boot from devices encrypted with LUKS2, but only if PBKDF2 is used. This means you cannot use Argon2 for a file system that contains the [/boot](#) directory. Note that even if PBKDF2 is used, some manual [grub2](#) configuration may be needed to boot from a LUKS2 device.

For more information on configuring device encryption with LUKS, use the [Help](#) button in the installer and refer to *Book “Security and Hardening Guide”, Chapter 13 “Storage encryption for hosted applications with cryptctl”*.

11.3 LVM configuration

This section explains specific steps to take when configuring LVM. If you need information about the Logical Volume Manager in general, refer to the *Book “Storage Administration Guide”, Chapter 5 “LVM configuration”, Section 5.1 “Understanding the logical volume manager”*.



Warning: Back up your data

Using LVM is sometimes associated with increased risk such as data loss. Risks also include application crashes, power failures, and faulty commands. Save your data before implementing LVM or reconfiguring volumes. Never work without a backup.

The YaST LVM configuration can be reached from the YaST Expert Partitioner (see [Section 11.1, “Using the Expert Partitioner”](#)) within the *Volume Management* item in the *System View* pane. The *Expert Partitioner* allows you to manage hard disks and partitions, as well as setting up RAID and LVM configurations.

11.3.1 Create physical volume

The first task is to create physical volumes that provide space to a volume group:

1. Select a hard disk from *Hard Disks*.

2. Change to the *Partitions* tab.
3. Click *Add* and enter the desired size of the PV on this disk.
4. Use *Do not format partition* and change the *File System ID* to *0x8E Linux LVM*. Do not mount this partition.
5. Repeat this procedure until you have defined all the desired physical volumes on the available disks.

11.3.2 Creating volume groups

If no volume group exists on your system, you must add one (see [Figure 11.3, “Creating a volume group”](#)). It is possible to create additional groups by clicking *Volume Management* in the *System View* pane, and then on *Add Volume Group*. One single volume group is usually sufficient.

1. Enter a name for the VG, for example, system.
2. Select the desired *Physical Extend Size*. This value defines the size of a physical block in the volume group. All the disk space in a volume group is handled in blocks of this size.
3. Add the prepared PVs to the VG by selecting the device and clicking *Add*. Selecting several devices is possible by holding **ctrl** while selecting the devices.
4. Select *Finish* to make the VG available to further configuration steps.




FIGURE 11.3: CREATING A VOLUME GROUP

If you have multiple volume groups defined and want to add or remove PVs, select the volume group in the *Volume Management* list and click *Resize*. In the following window, you can add PVs to or remove them from the selected volume group.

11.3.3 Configuring logical volumes

After the volume group has been filled with PVs, define the LVs which the operating system should use in the next dialog. Choose the current volume group and change to the *Logical Volumes* tab. *Add*, *Edit*, *Resize*, and *Delete* LVs as needed until all space in the volume group has been occupied. Assign at least one LV to each volume group.




FIGURE 11.4: LOGICAL VOLUME MANAGEMENT

Click *Add* and go through the wizard-like pop-up that opens:

1. Enter the name of the LV. For a partition that should be mounted to `/home`, a name like `HOME` could be used.
2. Select the type of the LV. It can be either *Normal Volume*, *Thin Pool*, or *Thin Volume*. Note that you need to create a thin pool first, which can store individual thin volumes. The big advantage of thin provisioning is that the total sum of all thin volumes stored in a thin pool can exceed the size of the pool itself.
3. Select the size and the number of stripes of the LV. If you have only one PV, selecting more than one stripe is not useful.
4. Choose the file system to use on the LV and the mount point.

By using stripes it is possible to distribute the data stream in the LV among several PVs (striping). However, striping a volume can only be done over different PVs, each providing at least the amount of space of the volume. The maximum number of stripes equals to the number of PVs, where Stripe "1" means "no striping". Striping only makes sense with PVs on different hard disks, otherwise performance will decrease.



Warning: Striping

YaST cannot verify your entries concerning striping at this point. Mistakes made here will show later when the LVM is implemented on disk.

If you have already configured LVM on your system, the existing logical volumes can also be used. Before continuing, assign appropriate mount points to these LVs. With *Finish*, return to the YaST *Expert Partitioner* and finish your work there.

11.4 Soft RAID

This section describes actions required to create and configure various types of RAID. In case you need background information about RAID, refer to *Book "Storage Administration Guide", Chapter 7 "Software RAID configuration", Section 7.1 "Understanding RAID levels"*.

11.4.1 Soft RAID configuration

The YaST *RAID* configuration can be reached from the YaST *Expert Partitioner*, described in [Section 11.1, "Using the Expert Partitioner"](#). This partitioning tool enables you to edit and delete existing partitions and create new ones to be used with soft RAID:

1. Select a hard disk from *Hard Disks*.
2. Change to the *Partitions* tab.
3. Click *Add* and enter the desired size of the raid partition on this disk.
4. Use *Do not Format the Partition* and change the *File System ID* to *0xFD Linux RAID*. Do not mount this partition.
5. Repeat this procedure until you have defined all the desired physical volumes on the available disks.

For RAID 0 and RAID 1, at least two partitions are needed—for RAID 1, usually exactly two and no more. If RAID 5 is used, at least three partitions are required, RAID 6 and RAID 10 require at least four partitions. It is recommended to use partitions of the same size only. The RAID partitions should be located on different hard disks to decrease the risk of losing data if one is defective (RAID 1 and 5) and to optimize the performance of RAID 0. After creating all the partitions to use with RAID, click *RAID > Add RAID* to start the RAID configuration.

In the next dialog, choose between RAID levels 0, 1, 5, 6 and 10. Then, select all partitions with either the “Linux RAID” or “Linux native” type that should be used by the RAID system. No swap or DOS partitions are shown.




FIGURE 11.5: RAID PARTITIONS

To add a previously unassigned partition to the selected RAID volume, first click the partition then *Add*. Assign all partitions reserved for RAID. Otherwise, the space on the partition remains unused. After assigning all partitions, click *Next* to select the available *RAID Options*.

In this last step, set the file system to use, encryption and the mount point for the RAID volume. After completing the configuration with *Finish*, see the `/dev/md0` device and others indicated with *RAID* in the *Expert Partitioner*.

11.4.2 Troubleshooting


Check the file `/proc/mdstat` to find out whether a RAID partition is damaged. If the system fails, shut down the machine and replace the defective hard disk with a new one partitioned the same way. Then restart your system and run `mdadm /dev/mdX --add /dev/sdX`. Replace 'X' with your particular device identifiers. This integrates the hard disk automatically into the RAID system and fully reconstructs it.

Note that although you can access all data during the rebuild, you may encounter some performance issues until the RAID has been fully rebuilt.

11.4.3 More information

Configuration instructions and more details for soft RAID can be found at:

- *Book “Storage Administration Guide”*
- <https://raid.wiki.kernel.org> 

Linux RAID mailing lists are available, such as <https://marc.info/?l=linux-raid> .

12 Remote installation

The installation of SUSE® Linux Enterprise Server can be performed entirely over the network. This chapter describes how to provide the required environment for booting, installing and controlling the installation via the network.

12.1 Overview

For a remote installation you need to consider how to boot, how to control the installation, and the source of the installation data. All available options can be combined with each other, if they are available for your hardware platform.

Boot method

Depending on the hardware, several options for booting a system exist. Common options are DVD, USB drive or PXE boot. For more information about your platform, refer to [Part I, “Installation preparation”](#).

To set up a server for booting via PXE, refer to [Chapter 18, Preparing network boot environment](#).

Data source

Most commonly, DVDs or USB drives are used as a source for installing SUSE Linux Enterprise Server. Alternatively, installation servers can be used. In this case, use the `install` boot parameter to specify the source. For details, refer to [Section 8.3.3, “Specifying the installation source”](#).

To use a network source for the installation, prepare a server as described in [Chapter 17, Setting up a network installation source](#).

Installation methods

Instead of using a keyboard and monitor directly attached to the target machine, the installation can be performed via SSH, VNC, or by using the serial console of a machine. This is described in the sections [Section 12.3, “Monitoring installation via VNC”](#), [Section 12.4, “Monitoring installation via SSH”](#) and [Section 12.5, “Installation via serial console”](#).

AutoYaST can be used to fully automate the installation process. For further information, refer to [Book “AutoYaST Guide”](#).

12.2 Scenarios for remote installation

This section introduces the most common installation scenarios for remote installations. For each scenario, carefully check the list of prerequisites and follow the procedure outlined for that scenario. If in need of detailed instructions for a particular step, follow the links provided for each one of them.

12.2.1 Installation from source media via VNC

This type of installation still requires some degree of physical access to the target system to boot for installation. The installation is controlled by a remote workstation using VNC to connect to the installation program. User interaction is required as with the manual installation in [Chapter 9, Installation steps](#).

For this type of installation, make sure that the following requirements are met.

- Target system with a working network connection.
- Controlling system with a working network connection and VNC viewer software or JavaScript-enabled browser (Firefox, Chromium, Internet Explorer, Opera, etc.).
- Installation DVD or USB flash drive.

To perform this kind of installation, proceed as follows:

1. Boot the target system using the installation medium (USB flash drive) of the SUSE Linux Enterprise Server media kit.
2. When the boot screen of the target system appears, use the boot parameters prompt to set the VNC options and static network configuration, if required. For information about boot parameters, see [Chapter 8, Boot parameters](#).

- a. Boot parameters for a static network configuration:

```
netdevice=NETDEVICE hostip=IP_ADDRESS netmask=NETMASK gateway=IP_GATEWAY vnc=1  
VNCPassword=PASSWORD
```

- b. Boot parameters for a dynamic (DHCP) network configuration:

```
vnc=1 VNCPassword=PASSWORD
```


3. The target system boots to a text-based environment and shows the network address and display number. VNC installations announce themselves over OpenSLP, provided the firewall settings are configured appropriately. They can be found using `slptool` as described in [Section 12.3.1, “Preparing for VNC installation”](#).
4. On the controlling workstation, open a VNC viewer or Web browser and connect to the target system using the provided network address and display number as described in [Section 12.3, “Monitoring installation via VNC”](#).
5. Perform the installation as described in [Chapter 9, Installation steps](#).

12.2.2 Network installation using VNC

This type of installation does not require a direct interaction with the target machine. The system is booted via PXE and the installation data is fetched from a server.

To perform this type of installation, make sure that the following requirements are met.

- At least one machine that can be used for installing a DHCP, NFS, HTTP, FTP, TFTP, or SMB server.
- Target system capable of PXE boot, networking, and Wake on LAN, plugged in and connected to the network.
- Controlling system with a working network connection and VNC viewer software or JavaScript-enabled browser (Firefox, Chromium, Microsoft Edge, Opera, etc.).

To perform this type of installation, proceed as follows.

1. Set up the server that contains the installation data. For details, see [Part IV, “Setting up an installation server”](#).
2. Set up a DHCP and TFTP server for the network. This is described in [Chapter 18, Preparing network boot environment](#). Add the required boot parameters to enable the VNC server.
3. Enable PXE boot in the target machine firmware. For more information, see [Section 18.4, “Preparing the target system for PXE boot”](#).
4. Initiate the boot process of the target system using Wake on LAN. This is described in [Section 18.5, “Using wake-on-LAN for remote wakeups”](#).
5. On the controlling workstation, open a VNC viewing application or Web browser and connect to the target system as described in [Section 12.3, “Monitoring installation via VNC”](#).

6. Perform the installation as described in [Chapter 9, Installation steps](#).

12.2.3 Installation from source media via SSH

This type of installation still requires some degree of physical access to the target system to boot for installation and to determine the IP address of the installation target. The installation itself is entirely controlled from a remote workstation using SSH to connect to the installer. User interaction is required as with the regular installation described in [Chapter 9, Installation steps](#).

For this type of installation, make sure that the following requirements are met.

- Target system with working network connection.
- Controlling system with working network connection and working SSH client software.
- Installation DVD or USB flash drive.

To perform this kind of installation, proceed as follows:

1. Set up the installation target and installation server as described in [Part IV, “Setting up an installation server”](#).
2. Boot the target system using the installation medium (USB flash drive) of the SUSE Linux Enterprise Server media kit.
3. When the boot screen of the target system appears, use the boot parameters prompt to set the SSH options and, if required, the static network configuration. For information about boot parameters, see [Chapter 8, Boot parameters](#).

- a. Boot parameters for a static network configuration:

```
netdevice=NETDEVICE hostip=IP_ADDRESS netmask=NETMASK gateway=IP_GATEWAY ssh=1  
ssh.password=PASSWORD
```

- b. Boot parameters for a dynamic (DHCP) network configuration:

```
ssh=1 ssh.password=PASSWORD
```

4. The target system boots to a text-based environment, giving the network address under which the graphical installation environment can be addressed by any SSH client.
5. On the controlling workstation, open a terminal window and connect to the target system as described in [Section 12.4.2, “Connecting to the installation program”](#).

6. Perform the installation as described in [Chapter 9, Installation steps](#).

12.2.4 Installation from network via SSH

This type of installation does not require a direct interaction with the target machine. The system is booted via PXE and the installation data is fetched from a server.

To perform this type of installation, make sure that the following requirements are met:

- At least one machine that can be used for installing a DHCP, NFS, HTTP, FTP, TFTP, or SMB server.
- Target system capable of PXE boot, networking, and Wake on LAN, plugged in and connected to the network.
- Controlling system with working network connection and SSH viewer software.

To perform this type of installation, proceed as follows.

1. Set up the server that contains the installation data. For details, see [Part IV, "Setting up an installation server"](#).
2. Set up a DHCP and TFTP server for the network. This is described in [Chapter 18, Preparing network boot environment](#). Add the required boot parameters to enable the SSH server.
3. Enable PXE boot in the target machine firmware. For more information, see [Section 18.4, "Preparing the target system for PXE boot"](#).
4. Initiate the boot process of the target system using Wake on LAN. This is described in [Section 18.5, "Using wake-on-LAN for remote wakeups"](#).
5. On the controlling workstation, open an SSH client software and connect to the target system as described in [Section 12.4, "Monitoring installation via SSH"](#).
6. Perform the installation as described in [Chapter 9, Installation steps](#).

12.3 Monitoring installation via VNC

Using a VNC viewer, you can remotely control the installation of SUSE Linux Enterprise Server from virtually any operating system. This section introduces the setup using a VNC viewer or a Web browser.

12.3.1 Preparing for VNC installation

To enable VNC on the installation target, specify the appropriate boot parameters at the initial boot for installation (see [Chapter 8, Boot parameters](#)). The target system boots into a text-based environment and waits for a VNC client to connect to the installation program.

The installation program announces the IP address and display number needed to connect for installation. If you have physical access to the target system, this information is provided right after the system booted for installation. Enter this data when your VNC client software prompts for it and provide your VNC password.

Because the installation target announces itself via OpenSLP, you can retrieve the address information of the installation target via an SLP browser. There is no need for physical access to the installation target provided your network setup and all machines support OpenSLP:

PROCEDURE 12.1: LOCATING VNC INSTALLATIONS VIA OPENSLP

1. Run `slptool findsrvtypes | grep vnc` to get a list of all services offering VNC. The VNC installation targets should be available under a service named `YaST.installation.suse`.
2. Run `slptool findsrvs YaST.installation.suse` to get a list of installations available. Use the IP address and the port (usually `5901`) provided with your VNC viewer.

12.3.2 Connecting to the installation program

There are two ways to connect to a VNC server (the installation target in this case). You can either start a VNC viewer or connect using a JavaScript-enabled Web browser.

Using VNC, you can install a Linux system from any other operating system, including other Linux distributions, Windows, or macOS.

On a Linux machine, make sure that the package `tightvnc` is installed. On a Windows machine, install the Windows port of this application (see <https://www.tightvnc.com/download.html> [↗](#)).

To connect to the installer running on the target machine, proceed as follows.

1. Start the VNC viewer.
2. Enter the IP address and display number of the installation target:

```
IP_ADDRESS:DISPLAY_NUMBER
```

This opens a window displaying the YaST screen as in a regular local installation.

Instead of a VNC viewer, you can use a JavaScript-enabled browser that has JavaScript support enabled to perform the installation.

Note that the browser VNC connection is not encrypted.

To perform a VNC installation, proceed as follows.

1. Launch the Web browser and enter the following at the address prompt:

```
http://IP_ADDRESS_OF_TARGET:5801
```

2. When prompted, enter the VNC password. This opens a window with the YaST screen as in a regular local installation.

12.4 Monitoring installation via SSH

Using an SSH client, you can perform the installation remotely via SSH.

12.4.1 Preparing for SSH installation

In addition to installing the required software package (OpenSSH for Linux and PuTTY for Windows), you need to specify the appropriate boot parameters to enable SSH for installation. See [Chapter 8, Boot parameters](#) for details. OpenSSH is installed by default on any SUSE Linux-based operating system.

12.4.2 Connecting to the installation program

After you have started the SSH installation, use this procedure to connect to the SSH session.

1. Retrieve the installation target's IP address. If you have physical access to the target machine, obtain the IP address that the installation routine provides from the console after the initial boot. Otherwise, obtain the IP address that has been assigned to the target machine in the DHCP server configuration.
2. Run the following command in the terminal:

```
ssh -X root@TARGET_IP_ADDRESS
```

Replace TARGET_IP_ADDRESS with the actual IP address of the installation target.

3. When prompted for a user name, enter root.
4. When prompted, enter the password that has been set with the SSH boot parameter. If the authentication is successful, you should see a command-line prompt for the installation target appear.
5. Enter **yast** to launch the installation program. This opens a window showing the YaST screen as described in [Chapter 9, Installation steps](#).

12.5 Installation via serial console

For this installation method, you need a computer connected by a *null modem* cable to the target machine where SUSE Linux Enterprise Server will be installed. Both machines must support the serial console. Certain firmware implementations are already configured to send the boot console output to a serial console. In this case, no additional configuration is required.

If the firmware does not use the serial console for the boot console output, set the following boot parameter for the installation: `console=TTY,BAUDRATE`. For further information, see *Book "Administration Guide", Chapter 18 "The boot loader GRUB 2", Section 18.2.5 "Editing menu entries during the boot procedure"* and [Chapter 8, Boot parameters](#).

`BAUDRATE` needs to be replaced by the baud rate for the interface. Valid values are 115200, 38400, or 9600. `TTY` needs to be replaced by the name of the interface. On most computers, there is one or more serial interfaces. Depending on the hardware, the names of the interfaces may vary:

- `ttyS0` for APM
- `ttyAMA0` for Server Base System Architecture (SBSA)
- `ttyPS0` for Xilinx

For the installation, you need a terminal program like **minicom** or **screen**. To initiate the serial connection, launch the screen program in a local console by entering the following command:

```
> screen /dev/ttyUSB0 115200
```

This means that screen listens to the first serial port with a baud rate of 115200. From this point on, the installation proceeds similarly to the text-based installation over this terminal.

13 Troubleshooting

This section covers several common installation problems and describes possible solutions.

13.1 Checking media

If you encounter any problems using the SUSE Linux Enterprise Server installation media, check its integrity. Boot from the media and choose *More > Check Installation Media* from the boot menu. A minimal system boots and lets you choose which device to check. Select the respective device and confirm with *OK* to perform the check.

On a running system, start YaST and choose *Software > Media Check*. Insert the medium and click *Start Check*. The integrity check may take time.

If errors are detected during the check, do not use this medium for installation. Media problems may, for example, occur when having burned the medium on DVD yourself. Burning the media at a low speed (4x) helps to avoid problems.

13.2 No bootable drive available

If your computer cannot boot from USB or DVD drive, you have several alternatives.

Using an external USB flash drive or DVD drive





Linux supports most existing USB flash drives and DVD drives. If the system has no USB flash drive or DVD drive, it is still possible that an external drive, connected through USB, FireWire, or SCSI, can be used to boot the system. Sometimes a firmware update may help if you encounter problems.

Network boot via PXE

If the machine lacks both a USB flash drive and DVD drive, but it has a working Ethernet connection, you can perform a network-based installation. See [Section 12.2.2, “Network installation using VNC”](#) and [Section 12.2.4, “Installation from network via SSH”](#) for details.

USB flash drive

You can use a USB flash drive if the machine lacks a DVD drive and a network connection. For details, see:

-  [Section 2.4, “Booting the system”](#) 
-  [Section 3.4, “Booting the system”](#) 

13.3 Booting from installation media fails

The machine may fail to boot from the installation media due to an incorrect boot sequence setting in BIOS. The USB flash drive or DVD drive must be set as the first boot device in the BIOS boot sequence.

PROCEDURE 13.1: CHANGING THE BIOS BOOT SEQUENCE

1. Enter the BIOS using the proper key shown by the boot routines and wait for the BIOS screen to appear.
2. To change the boot sequence in an AWARD BIOS, look for the *BIOS FEATURES SETUP* entry. Other manufacturers may have a different name for this, such as *ADVANCED CMOS SETUP*. When you have found the entry, select it and confirm with **Enter**.
3. Look for a subentry called *BOOT SEQUENCE* or *BOOT ORDER*. Change the settings by pressing **Page ↑** or **Page ↓** until the USB flash drive or DVD drive is listed first.
4. Exit the BIOS setup screen by pressing **Esc**. To save the changes, select *SAVE & EXIT SETUP*, or press **F10**. To save the modified settings, press **Y**.

PROCEDURE 13.2: CHANGING THE BOOT SEQUENCE IN AN SCSI BIOS (ADAPTEC HOST ADAPTER)

1. Open the setup by pressing **Ctrl – A**.
2. Select *Disk Utilities*. The connected hardware components are now displayed. Make note of the SCSI ID of your USB flash drive or DVD drive.
3. Exit the menu with **Esc**.
4. Open *Configure Adapter Settings*. Under *Additional Options*, select *Boot Device Options* and press **Enter**.
5. Enter the ID of the USB flash drive or DVD drive and press **Enter** again.
6. Press **Esc** twice to return to the start screen of the SCSI BIOS.

7. Exit this screen and confirm with *Yes* to boot the computer.

Regardless of what language and keyboard layout the installed system will be using, most BIOS configurations use the US keyboard layout as shown below.



FIGURE 13.1: US KEYBOARD LAYOUT

13.4 Boot failure

Some hardware types, mainly very old or very recent ones, fail to boot. Reasons can be missing support for hardware in the installation kernel or drivers causing problems on some specific hardware.

If installation fails using the standard *Installation* mode, try the following.

1. With the installation media still in the drive, reboot the machine with **Ctrl – Alt – Del** or using the hardware reset button.
2. When the boot screen appears, press **F5**, use the arrow keyboard keys to navigate to *No ACPI*, and press **Enter** to boot and initiate the installation process. This option disables the support for ACPI power management techniques.
3. Proceed with the installation as described in [Chapter 9, Installation steps](#).

If this fails, proceed as above, but choose *Safe Settings* instead. This option disables ACPI and DMA support. This option works with most hardware.

If both options fail, use the boot parameters prompt to specify the kernel parameters to enable support for the hardware in use. For more information about the parameters available as boot parameters, refer to the kernel documentation located in [/usr/src/linux/Documentation/kernel-parameters.txt](#).



Tip: Obtaining kernel documentation

Install the kernel-source package to view the kernel documentation.

There are other ACPI-related kernel parameters that can be entered at the boot prompt prior to booting for installation:

acpi=off

This parameter disables the complete ACPI subsystem on your computer. This may be useful if your computer cannot handle ACPI or if you think ACPI in your computer causes trouble.

acpi=force

Always enable ACPI even if your computer has a BIOS released before 2000. This parameter also enables ACPI if it is set in addition to acpi=off.

acpi=noirq

Do not use ACPI for IRQ routing.

acpi=ht

Run only enough ACPI to enable hyper-threading.

acpi=strict

Be less tolerant of platforms that are not strictly ACPI-compliant.

pci=noacpi

Disable PCI IRQ routing of the new ACPI system.

pnpacpi=off

Enable this option to avoid issues caused by incorrectly configured device resources in BIOS.

notsc

Disable the time stamp counter. This option can be used to work around timing problems on your systems. It is a recent feature, so if you see regressions on your machine, especially time related or even total hangs, this option is worth a try.

nohz=off

Disable the nohz feature. If your machine hangs, enabling this option may help.

When you have determined the right parameter combination, YaST automatically writes them to the boot loader configuration to make sure that the system boots properly next time.

If errors occur when the kernel is loaded or during the installation, select *Memory Test* in the boot menu to check the memory. If *Memory Test* returns an error, this usually indicates a hardware error.

13.5 Graphical installer fails to start

The machine boots into the installation interface, and the graphical installer does not start when you select *Installation*.

There are several ways to deal with this situation.

- Select another screen resolution for the installation dialogs.
- Select *Text Mode* for installation.
- Perform a remote installation via VNC using the graphical installer.

PROCEDURE 13.3: CHANGE SCREEN RESOLUTION FOR INSTALLATION

1. Boot for installation.
2. Press **F3** to open a menu from which to select a lower resolution for installation purposes.
3. Select *Installation* and proceed with the installation as described in [Chapter 9, Installation steps](#).

PROCEDURE 13.4: INSTALLATION IN TEXT MODE

1. Boot for installation.
2. Press **F3** and select *Text Mode*.
3. Select *Installation* and proceed with the installation as described in [Chapter 9, Installation steps](#).

PROCEDURE 13.5: VNC INSTALLATION

1. Boot for installation.
2. Enter the following text at the boot parameters prompt:

```
vnc=1 vncpassword=SOME_PASSWORD
```

Replace SOME_PASSWORD with the password to use for VNC installation.

3. Select *Installation* then press **Enter** to start the installation.

Instead of starting right into the graphical installation routine, the system continues to run in a text mode. The system then halts, displaying a message containing the IP address and port number at which the installer can be reached via a browser interface or a VNC viewer application.

4. When using a browser to access the installer, launch the browser and enter the address information provided by the installation routines on the future SUSE Linux Enterprise Server machine and press **Enter** :

```
http://IP_ADDRESS_OF_MACHINE:5801
```

A dialog opens in the browser window prompting you for the VNC password. Enter it and proceed with the installation as described in [Chapter 9, Installation steps](#).



Important: Cross-platform support

Installation via VNC works with any browser under any operating system, provided Java support is enabled.

Provide the IP address and password to your VNC viewer when prompted. A window opens, displaying the installation dialogs. Proceed with the installation as usual.

13.6 Only minimal boot screen is displayed

You inserted the medium into the drive, the BIOS routines are finished, and the system launches a minimal, text-based interface. This may happen on any machine that does not have sufficient graphics memory for rendering a graphical boot screen.

Although the text boot screen looks minimal, it provides nearly the same functionality as the graphical one.

Boot options

Unlike the graphical interface, the different boot parameters cannot be selected using the cursor keys of your keyboard. The boot menu of the text-mode boot screen provides keywords that can be entered at the boot prompt. These keywords match the options in the graphical version. Enter your choice and press **Enter** to launch the boot process.

Custom boot options

After selecting a boot parameter, enter the appropriate keyword at the boot prompt or enter some custom boot parameters as described in [Section 13.4, “Boot failure”](#). To launch the installation process, press **Enter** .

Screen resolutions

Use the function keys (**F1** ... **F12**) to determine the screen resolution for installation. If you need to boot in text mode, choose **F3** .

13.7 Log files

For more information about log files created during installation, see *Book “Administration Guide”, Chapter 47 “Gathering system information for support”, Section 47.5 “Gathering information during the installation”*.

III Customizing installation images

- 14 Prepare a disk for cloning with the system cleanup tool **197**
- 15 Customizing installation images with mksusecd **198**
- 16 Customizing installation images manually **202**

14 Prepare a disk for cloning with the system cleanup tool

The `clone-master-clean-up` tool that ships with SUSE Linux Enterprise Server makes it possible to remove data from the disk that you do not want to include in a clone. This chapter describes how to use the tool.

14.1 Cleaning up unique system identifiers



Warning: Do not use the tool on a production system

As the cleanup tool removes essential system configuration data, it is not recommended to use it on a system that is used in production. Run the tool on the cloned image instead.

The `clone-master-clean-up` tool removes the following data:

- swap files
- Zypper repositories
- SSH host and client keys
- temporary directories, like `/tmp/*`
- Postfix data
- HANA firewall script
- systemd journal

1. To install `clone-master-clean-up`, run the following command:

```
> sudo zypper install clone-master-clean-up
```

2. Configure the tool by editing the `/etc/sysconfig/clone-master-clean-up` file. Here, you can specify which specific data the tool should remove.
3. Run the script to perform a cleanup:

```
> sudo clone-master-clean-up
```

15 Customizing installation images with mksusecd

mksusecd is a useful tool for creating a customized installation image. Use it to modify the regular SUSE Linux Enterprise installation images, add or remove files, create a minimal network installation image, customize boot options or software repositories, and create a minimal boot image as an alternative to booting a system from a PXE server.

15.1 Installing mksusecd

In SLE 15, **mksusecd** is in the Development Tools module. If this module is not enabled, you must enable it first. Find the exact module name and **SUSEConnect** activation command with **zypper**:

```
> zypper search-packages mksusecd
Following packages were found in following modules:

Package                Module or Repository
-----
mksusecd               Development Tools Module (sle-module-development-tools/15.4/x86_64)
                        SUSEConnect --product sle-module-development-tools/15.4/x86_64

To activate the respective module or product, use SUSEConnect --product.
Use SUSEConnect --help for more details.
```

Enable the module with SUSEConnect:

```
> sudo SUSEConnect --product sle-module-development-tools/15.4/x86_64
```

Install **mksusecd**:

```
> sudo zypper in mksusecd
```

Run **mksusecd --help** to see a complete list of commands.

After you create your custom image, either burn it to a CD/DVD medium using your preferred disk-writing program, or create a bootable USB flash drive using the **dd** command. Make sure the device is not mounted, then run the following command:

```
# dd if=myinstaller.iso of=/dev/SDB bs=4M
```


Then your new bootable device is ready to use.

15.2 Creating a minimal boot image

Use **mksusecd** to create a minimal boot image to start client machines from a CD/DVD or USB flash drive, instead of starting them from a PXE boot server. The minimal boot image launches the kernel and initrd, and then the remaining installation files are fetched from a local NFS server (see [Section 17.1, “Setting up an installation server using YaST”](#)).

Run the following command to create the minimal ISO image:

```
> sudo mksusecd --create min-install.iso \  
--net=nfs://192.168.1.1:/srv/install/ARCH/OS_VERSION/SP_VERSION/cd1 \  
/srv/tftpboot/EFI/ARCH/boot
```

Replace the NFS server address with your own. Replace ARCH with the directory corresponding to the target system architecture. Also replace OS_version and SP_VERSION (service pack) according to your paths in [Section 17.1, “Setting up an installation server using YaST”](#).

15.3 Setting default kernel boot parameters

Rather than waiting for a boot prompt to enter your custom kernel boot parameters, configure them in a custom **mksusecd** image:

```
> sudo mksusecd --create install.iso \  
--boot "textmode=1 splash=silent mitigations=auto"
```

Verify that your custom parameters load correctly after start-up by querying /proc:

```
> cat /proc/cmdline
```

15.4 Customizing modules, extensions, and repositories

SUSE Linux Enterprise 15 supports Modules (not to be confused with kernel modules) and Extensions for different product components. These are add-ons to the default Basesystem, such as Development Tools, Desktop Applications, and SUSE Linux Enterprise Live Patching. For more information refer to the Modules and Extensions Quick Start guide.

With **mksusecd** you can create an installation image containing all additional Modules and Extensions you want. Start by querying existing images, like this example for SUSE Linux Enterprise 15 SP6:

```
> sudo mksusecd --list-repos SLE-15-SP6-Full-ARCH-GM-media1.iso
Repositories:
  Basesystem-Module [15.6-0]
  SUSE-CAP-Tools-Module [15.6-0]
  Containers-Module [15.6-0]
  Desktop-Applications-Module [15.6-0]
  Development-Tools-Module [15.6-0]
  HPC-Module [15.6-0]
  Legacy-Module [15.6-0]
  Live-Patching [15.6-0]
  Public-Cloud-Module [15.6-0]
  Python2-Module [15.6-0]
  SAP-Applications-Module [15.6-0]
  Server-Applications-Module [15.6-0]
  Transactional-Server-Module [15.6-0]
  Web-Scripting-Module [15.6-0]
  SLEHA15-SP6 [15.6-0]
  SLE-15-SP6-HPC [15.6-0]
  SLED15-SP6 [15.6-0]
  SLES15-SP6 [15.6-0]
  SLE-15-SP6-SAP [15.6-0]
  SLEWE15-SP6 [15.6-0]
  [...]
```

Create a new installation image that is built from the Modules, Extensions, and repositories that you select, and automatically enable them:

```
> sudo mksusecd --create myinstaller.iso --enable-repos auto \
--include-repos Basesystem-Module,Desktop-Applications-Module \
SLE-15-SP6-Full-ARCH-GM-media1.iso
```

This example creates an image for installation from the internet. To create an image for offline installation, additionally add the repository of the base product, for example SLES15-SP6 for SUSE Linux Enterprise Server.

```
> sudo mksusecd --create myinstaller.iso --enable-repos auto \
--include-repos SLES15-SP6,Basesystem-Module,Desktop-Applications-Module \
SLE-15-SP6-Full-ARCH-GM-media1.iso
```

Replace --enable-repos auto with --enable-repos ask to have the installer present a dialog for choosing modules.



Note: AutoYaST control file

When using the `--enable-repos` option, `mksusecd` adds an `add_on_products.xml` file for use with AutoYaST to the new image. Modules in this file do not need to be listed in the in the AutoYaST control file.

15.5 Creating a minimal netinstall ISO

To create a minimal installation image to launch a network installation, use the `--nano` option:


```
> sudo mksusecd --create netinstall.iso \  
--nano SLE-15-SP6-Online-ARCH-GM-media1.iso
```

15.6 Change default repository

To set a different repository, such as your own local repository, use the `--net` option:

```
> sudo mksusecd --create localinstall.iso \  
--net "https://example.com/local" SLE-15-SP6-Online-ARCH-GM-media1.iso
```

16 Customizing installation images manually

You can customize the standard SUSE Linux Enterprise installation images by editing a file in the installation ISO image, `media.1/products`. Add modules and extensions to create a single customized installation image. Then copy your custom image to a CD, DVD, or USB flash drive to create a customized bootable installation medium. See [the SUSE Best Practices paper on *How to Create a Custom Installation Medium for SUSE Linux Enterprise 15*](https://documentation.suse.com/sbp/all/single-html/SBP-SLE15-Custom-Installation-Medium/) (<https://documentation.suse.com/sbp/all/single-html/SBP-SLE15-Custom-Installation-Medium/>)  for complete instructions.

IV Setting up an installation server

- 17 Setting up a network installation source **204**
- 18 Preparing network boot environment **214**
- 19 Setting up a UEFI HTTP Boot server **226**
- 20 Deploying customized preinstallations **235**

17 Setting up a network installation source

This chapter describes how to create a server that provides the data required for installing SUSE Linux Enterprise Server over the network.

Depending on the operating system of the machine used as the network installation source for SUSE Linux Enterprise Server, there are several options for the server configuration. The easiest way to set up an installation server is to use YaST.



Tip: Installation server operating system

You can even use a Microsoft Windows machine as the installation server for your Linux deployment. See [Section 17.5, “Managing an SMB repository”](#) for details.

17.1 Setting up an installation server using YaST

YaST offers a graphical tool for creating network repositories. It supports HTTP, FTP, and NFS network installation servers.

1. Log in to the machine that should act as installation server.
2. Install the package `yast2-instserver`:

```
> sudo zypper in yast2-instserver
```

3. Start *YaST* > *Miscellaneous* > *Installation Server*.
4. Select the repository type (HTTP, FTP, or NFS). The selected service is started automatically every time the system starts. If a service of the selected type is already running on your system and you want to configure it manually for the server, deactivate the automatic configuration of the server service with *Do Not Configure Any Network Services*. In both cases, define the directory in which the installation data should be made available on the server.
5. Configure the required repository type. This step relates to the automatic configuration of server services. It is skipped when automatic configuration is deactivated.

Define an alias for the root directory of the FTP or HTTP server on which the installation data should be found. The repository will later be located under `ftp://Server-IP/Alias/Name` (FTP) or under `http://Server-IP/Alias/Name` (HTTP). *Name* stands

for the name of the repository, which is defined in the following step. If you selected NFS in the previous step, define wild cards and export options. The NFS server will be accessible under `nfs://Server-IP/Name`. Details of NFS and exports can be found in *Book "Storage Administration Guide", Chapter 19 "Sharing file systems with NFS"*.



Tip: Firewall settings

Make sure that the firewall settings of your server system allow traffic on ports for HTTP, NFS, and FTP. If they currently do not, enable *Open Port in Firewall* or check *Firewall Details* first.

6. Configure the repository. Before the installation media are copied to their destination, define the name of the repository (ideally, an easily remembered abbreviation of the product and version). YaST allows providing ISO images of the media instead of copies of the installation DVDs. If you want this, activate the relevant check box and specify the directory path under which the ISO files can be found locally. Depending on the product to distribute using this installation server, it might be necessary to add media, such as service pack DVDs as extra repositories. To announce the installation server on the network via OpenSLP, activate the appropriate option.



Tip: Announcing the repository

Consider announcing your repository via OpenSLP if your network setup supports this option. This saves you from entering the network installation path on every target machine. The target systems are booted using the SLP boot parameter and find the network repository without any further configuration. For details on this option, refer to [Chapter 8, Boot parameters](#).

7. Configuring extra repositories. YaST follows a specific naming convention to configure add-on CD or service pack CD repositories. Configuration is accepted only if the repository name of the add-on CDs starts with the repository name of the installation media. In other words, if you chose SLES12SP1 as repository name for DVD1 then you should select SLES12SP1addon as repository name for DVD2.
8. Upload the installation data. The most lengthy step in configuring an installation server is copying the actual installation media. Insert the media in the sequence requested by YaST and wait for the copying procedure to end. When the sources have been fully copied, return to the overview of existing repositories and close the configuration by selecting *Finish*.

Your installation server is now fully configured and ready for service. It is automatically started every time the system is started. No further intervention is required. You only need to configure and start this service correctly manually if you deactivated the automatic configuration of the selected network service with YaST as an initial step.

To deactivate a repository, select the repository to remove then select *Delete*. The installation data are removed from the system. To deactivate the network service, use the respective YaST module.

If your installation server needs to provide the installation data for more than one product of the product version, start the YaST installation server module. Then select *Add* in the overview of existing repositories to configure the new repository.



Warning: YaST installation server will conflict with RMT server

Configuring a server to be an installation server with YaST automatically installs and configures the Apache Web server, listening on port 80.

However, configuring a machine to be an RMT server (Repository Mirroring Tool) automatically installs the NGINX Web server and configures it to listen on port 80.

Do not try to enable both these functions on the same server. It is not possible for a single server to host both simultaneously.

17.2 Setting up an NFS repository manually

Setting up an NFS source for installation is done in two main steps. First, create the directory structure holding the installation data and copy the installation media over to this structure. Second, export the directory holding the installation data to the network.

To create a directory to hold the installation data, proceed as follows:

1. Log in as root.
2. Create a directory that will hold all installation data and change into this directory. For example:

```
# mkdir -p /srv/install/PRODUCT/PRODUCTVERSION
# cd /srv/install/PRODUCT/PRODUCTVERSION
```


Replace *PRODUCT* with an abbreviation of the product name and *PRODUCTVERSION* with a string that contains the product name and version (for example, */srv/install/SLES/15.1*).

3. For each installation medium contained in the media kit, execute the following commands:

- a. Copy the entire content of the installation medium into the installation server directory:

```
# cp -a /media/PATH_TO_YOUR_MEDIA_DRIVE .
```

Replace *PATH_TO_YOUR_MEDIA_DRIVE* with the actual mount point of the installation medium.

- b. Rename the directory to the medium number:

```
# mv PATH_TO_YOUR_MEDIA_DRIVE DVDX
```

Replace *X* with the actual number of the installation medium.

On SUSE Linux Enterprise Server, you can export the repository with NFS using YaST. Proceed as follows:

1. Log in as *root*.
2. Start *YaST > Network Services > NFS Server*.
3. Select *Start* and *Open Port in Firewall* and click *Next*.
4. Select *Add Directory* and browse for the directory containing the installation sources, in this case, *PRODUCTVERSION*.
5. Select *Add Host* and enter the host names of the machines to which to export the installation data. Instead of specifying host names here, you could also use wild cards, ranges of network addresses, or the domain name of your network. Enter the appropriate export options or leave the default, which works fine in most setups. For more information about the syntax used in exporting NFS shares, read the *exports* man page.
6. Click *Finish*. The NFS server holding the SUSE Linux Enterprise Server repository is automatically started and integrated into the boot process.

To export the repository manually via NFS instead of using the YaST NFS Server module, proceed as follows:

1. Log in as root.
2. Open the file /etc/exports and enter the following line:

```
/PRODUCTVERSION *(ro,root_squash,sync)
```

This exports the directory /PRODUCTVERSION to any host that is part of this network or to any host that can connect to this server. To limit the access to this server, use netmasks or domain names instead of the general wild card *. Refer to the export man page for details. Save and exit this configuration file.

3. To add the NFS service to the list of servers started during system boot, execute the following commands:

```
# systemctl enable nfsserver
```

4. Start the NFS server with **systemctl start nfsserver**. If you need to change the configuration of your NFS server later, modify the configuration file and restart the NFS daemon with **systemctl restart nfsserver**.

Announcing the NFS server via OpenSLP makes its address known to all clients in your network.

1. Log in as root.
2. Create the /etc/slp.reg.d/install.suse.nfs.reg configuration file with the following lines:

```
# Register the NFS Installation Server
service:install.suse:nfs://$HOSTNAME/PATH_TO_REPOSITORY/DVD1,en,65535
description=NFS Repository
```

Replace PATH_TO_REPOSITORY with the actual path to the installation source on your server.

3. Start the OpenSLP daemon with **systemctl start slpd**.

For more information about OpenSLP, refer to the package documentation located under /usr/share/doc/packages/openslp/ or refer to *Book "Administration Guide", Chapter 41 "SLP"*. For more information about NFS, refer to *Book "Storage Administration Guide", Chapter 19 "Sharing file systems with NFS"*.

17.3 Setting up an FTP repository manually

Creating an FTP repository is very similar to creating an NFS repository. An FTP repository can be announced over the network using OpenSLP as well.

1. Create a directory holding the installation sources as described in [Section 17.2, “Setting up an NFS repository manually”](#).
2. Configure the FTP server to distribute the contents of your installation directory:

- a. Log in as `root` and install the package `vsftpd` using the YaST software management.
- b. Enter the FTP server root directory:

```
# cd /srv/ftp
```

- c. Create a subdirectory holding the installation sources in the FTP root directory:

```
# mkdir REPOSITORY
```

Replace `REPOSITORY` with the product name.

- d. Mount the contents of the installation repository into the change root environment of the FTP server:

```
# mount --bind PATH_TO_REPOSITORY /srv/ftp/REPOSITORY
```

Replace `PATH_TO_REPOSITORY` and `REPOSITORY` with values matching your setup. If you need to make this permanent, add it to `/etc/fstab`.

- e. Start `vsftpd` with `vsftpd`.

3. Announce the repository via OpenSLP, if this is supported by your network setup:

- a. Create the `/etc/slp.reg.d/install.suse.ftp.reg` configuration file with the following lines:

```
# Register the FTP Installation Server
service:install.suse:ftp://$HOSTNAME/REPOSITORY/DVD1,en,65535
description=FTP Repository
```

Replace `REPOSITORY` with the actual name of the repository directory on your server. The `service:` line should be entered as one continuous line.

- b. Start the OpenSLP daemon with `systemctl start slpd`.



Tip: Configuring an FTP server with YaST

If you prefer to use YaST rather than manually configuring the FTP installation server, refer to *Book "Administration Guide", Chapter 43 "Setting up an FTP server with YaST"*.

17.4 Setting up an HTTP repository manually

Creating an HTTP repository is very similar to creating an NFS repository. An HTTP repository can be announced over the network using OpenSLP as well.

1. Create a directory holding the installation sources as described in [Section 17.2, "Setting up an NFS repository manually"](#).
2. Configure the HTTP server to distribute the contents of your installation directory:
 - a. Install the Web server Apache as described in *Book "Administration Guide", Chapter 42 "The Apache HTTP server", Section 42.1.2 "Installation"*.
 - b. Enter the root directory of the HTTP server (`/srv/www/htdocs`) and create the sub-directory that will hold the installation sources:

```
# mkdir REPOSITORY
```

Replace `REPOSITORY` with the product name.

- c. Create a symbolic link from the location of the installation sources to the root directory of the Web server (`/srv/www/htdocs`):

```
# ln -s /PATH_TO_REPOSITORY/srv/www/htdocs/REPOSITORY
```

- d. Modify the configuration file of the HTTP server (`/etc/apache2/default-server.conf`) to make it follow symbolic links. Replace the following line:

```
Options None
```

with

```
Options Indexes FollowSymLinks
```

- e. Reload the HTTP server configuration using `systemctl reload apache2`.

3. Announce the repository via OpenSLP, if this is supported by your network setup:

- a. Create the `/etc/slp.reg.d/install.suse.http.reg` configuration file with the following lines:

```
# Register the HTTP Installation Server
service:install.suse:http://$HOSTNAME/REPOSITORY/DVD1/,en,65535
description=HTTP Repository
```

Replace `REPOSITORY` with the actual path to the repository on your server. The `service:` line should be entered as one continuous line.

- b. Start the OpenSLP daemon using `systemctl start slpd`.

17.5 Managing an SMB repository

Using SMB, you can import the installation sources from a Microsoft Windows server and start your Linux deployment even with no Linux machine around.

To set up an exported Windows Share holding your SUSE Linux Enterprise Server repository, proceed as follows:

1. Log in to your Windows machine.
2. Create a new directory that will hold the entire installation tree and name it `INSTALL`, for example.
3. Export this share according to the procedure outlined in your Windows documentation.
4. Enter this share and create a subdirectory, called `PRODUCT`. Replace `PRODUCT` with the actual product name.
5. Enter the `INSTALL/PRODUCT` directory and copy each medium to a separate directory, such as `DVD1` and `DVD2`.

To use an SMB mounted share as a repository, proceed as follows:

1. Boot the installation target.
2. Select *Installation*.
3. Press **F4** for a selection of the repository.

4. Choose SMB and enter the Windows machine's name or IP address, the share name (INSTALL/PRODUCT/DVD1, in this example), user name, and password. The syntax looks like this:

```
smb://workdomain;user:password@server/INSTALL/DVD1
```

After you press **Enter**, YaST starts and you can perform the installation.

17.6 Using ISO images of the installation media on the server

Instead of copying physical media into your server directory manually, you can also mount the ISO images of the installation media into your installation server and use them as a repository. To set up an HTTP, NFS or FTP server that uses ISO images instead of media copies, proceed as follows:

1. Download the ISO images and save them to the machine to use as the installation server.
2. Log in as root.
3. Choose and create an appropriate location for the installation data, as described in *Section 17.2, "Setting up an NFS repository manually"*, *Section 17.3, "Setting up an FTP repository manually"*, or *Section 17.4, "Setting up an HTTP repository manually"*.
4. Create subdirectories for each installation medium.
5. To mount and unpack each ISO image to the final location, issue the following command:

```
# mount -o loop PATH_TO_ISO PATH_TO_REPOSITORY/PRODUCT/MEDIUMX
```

Replace PATH_TO_ISO with the path to your local copy of the ISO image. Replace PATH_TO_REPOSITORY with the source directory of your server. Replace PRODUCT with the product name and replace MEDIUMX with the type (CD or DVD) and number of media you are using.

6. Repeat the previous step to mount all ISO images needed for your product.
7. Start your installation server as usual, as described in *Section 17.2, "Setting up an NFS repository manually"*, *Section 17.3, "Setting up an FTP repository manually"*, or *Section 17.4, "Setting up an HTTP repository manually"*.

To automatically mount the ISO images at boot time, add the respective mount entries to /etc/fstab. An entry according to the previous example would look like the following:

```
PATH_TO_ISO PATH_TO_REPOSITORY/PRODUCTMEDIUM auto loop
```

18 Preparing network boot environment

This chapter describes how to configure a DHCP and a TFTP server that provide the required infrastructure for booting with PXE.

SUSE® Linux Enterprise Server can be installed via a Preboot Execution Environment (PXE). The client hardware needs to support booting via PXE. The network needs to provide a DHCP server and a TFTP server providing the required data to the clients. This chapter guides you through setting up the required servers.

PXE only boots a kernel and initrd. This can be used to boot into an installation environment or into live systems. To set up the installation sources, see [Chapter 17, Setting up a network installation source](#).

This section covers the configuration tasks needed in complex boot scenarios. It contains ready-to-apply configuration examples for DHCP, PXE boot, TFTP, and Wake on LAN.

The examples assume that the DHCP, TFTP and NFS server reside on the same machine with the IP 192.168.1.1. All services can reside on different machines without any problems. Make sure to change the IP addresses as required.

18.1 Setting up a DHCP server

A DHCP server provides both dynamic ([Section 18.1.1, “Dynamic address assignment”](#)) and static IP address assignments ([Section 18.1.2, “Assigning static IP addresses”](#)) to your network clients. It advertises servers, routes, and domains. For TFTP servers, DHCP also provides the kernel and initrd files. Which files are loaded depends on the architecture of the target machine, and whether legacy BIOS or UEFI boot is used. The clients transmit their architecture type in their DHCP requests. Based on this information, the DHCP server decides which files the client must download for booting.



Warning: PXE and AutoYaST installation failure

Starting with SUSE Linux Enterprise 15.0, there are special conditions that cause PXE boot and AutoYaST installations to fail. See [Section 18.1.3, “PXE and AutoYaST installation failures”](#) for more information and the solution.

18.1.1 Dynamic address assignment

The following example shows how to set up a DHCP server that dynamically assigns IP addresses to clients, and advertises servers, routers, domains, and boot files.

1. Log in as root to the machine hosting the DHCP server.
2. Enable the DHCP server by executing **`systemctl enable dhcpd`**.
3. Append the following lines to a subnet configuration of your DHCP server's configuration file located under `/etc/dhcpd.conf`:

```
# The following lines are optional
option domain-name "my.lab";
option domain-name-servers 192.168.1.1;
option routers 192.168.1.1;
option ntp-servers 192.168.1.1;
ddns-update-style none;
default-lease-time 3600;

# The following lines are required
option arch code 93 = unsigned integer 16; # RFC4578
subnet 192.168.1.0 netmask 255.255.255.0 {
    next-server 192.168.1.1;
    range 192.168.1.100 192.168.1.199;
    default-lease-time 3600;
    max-lease-time 3600;
    if option arch = 00:07 or option arch = 00:09 {
        filename "/EFI/x86/grub.efi";
    }
    else if option arch = 00:0b {
        filename "/EFI/aarch64/bootaa64.efi";
    }
    else {
        filename "/BIOS/x86/pxelinux.0";
    }
}
```

This configuration example uses the subnet 192.168.1.0/24 with the DHCP, DNS and gateway on the server with the IP 192.168.1.1. Make sure that all IP addresses are changed according to your network layout. For more information about the options available in `dhcpd.conf`, refer to the `dhcpd.conf` manual page.

4. Restart the DHCP server by executing **`systemctl restart dhcpd`**.

18.1.2 Assigning static IP addresses

A DHCP server may also assign static IP addresses and host names to network clients. One use case is assigning static addresses to servers. Another use case is restricting which clients may join the network to those with assigned static IP addresses, and providing no dynamic address pools.

Modify the above DHCP configuration according to the following example:

```
group {
  host test {
    hardware ethernet MAC_ADDRESS;
    fixed-address IP_ADDRESS;
  }
}
```

The host statement assigns a host name to the installation target. To bind the host name and IP address to a specific host, you must specify the client's hardware (MAC) address. Replace all the variables used in this example with the actual values that match your environment, then save your changes and restart the DHCP server.

18.1.3 PXE and AutoYaST installation failures

Starting with SUSE Linux Enterprise 15.0 and ISC DHCP 4.3.x, there are special circumstances that cause PXE boot and AutoYaST installations to fail. If your DHCP server does not have a pool of available dynamic IP addresses, but allows only pre-defined static addresses per client, and the clients send RFC 4361 client identifiers, then PXE/AutoYaST installations will not work. (Allowing only addresses assigned to specific network clients, and providing no dynamic address pools, prevents random machines from joining the network.)

When a new system starts in PXE, it sends a request to the DHCP server and identifies itself using a client identifier constructed from the hardware type plus the MAC address of the network interface. This is an RFC 2132 `client-id`. The DHCP server then offers the assigned IP address. Next, the installation kernel is loaded, and sends another DHCP request, but this `client-id` is different, and is sent in RFC 4361 format. The DHCP server will not recognize this as the same client, and will look for a free dynamic IP address, which is not available, and the installation stops.

The solution is to configure clients to send RFC 2132 client IDs. To send an RFC 2132 `client-id` during the installation, use `linuxrc` to pass the following `ifcfg` command:

```
ifcfg=eth0=dhcp,DHCLIENT_CLIENT_ID=01:03:52:54:00:02:c2:67,
DHCLIENT6_CLIENT_ID=00:03:52:54:00:02:c2:67
```

The traditionally-used RFC 2132 DHCPv4 `client-id` on Ethernet is constructed from the hardware type (`01` for Ethernet) and followed by the hardware address (the MAC address), for example:

```
01:52:54:00:02:c2:67
```

The RFC 4361 DHCPv4 `client-id` attempts to correct the problem of identifying a machine that has more than one network interface. The new DHCPv4 `client-id` has the same format as the DHCPv6 `client-id`. It starts with the `0xff` prefix, instead of the hardware type, followed by the DHCPv6 IAID (the interface-address association ID that describes the interface on the machine), followed by the DHCPv6 Unique Identifier (DUID), which uniquely identifies the machine.

Using the above hardware type-based and hardware address-based DUID, the new RFC 4361 DHCPv4 `client-id` would be:

- Using the last bytes of the MAC address as the IAID:
`ff:00:02:c2:67:00:01:xx:xx:xx:xx:52:54:00:02:c2:67`
- When the IAID is a simple incremented number:
`ff:00:00:00:01:00:01:xx:xx:xx:xx:52:54:00:02:c2:67`

The `xx:xx:xx:xx` field in the DUID-Link-Layer Timestamp (DUID-LLT) is a creation time stamp. A DUID-Link-Layer (DUID-LL) (`00:03:00:01:$MAC`) does not have a time stamp.

For more information on using `linuxrc`, see the AutoYaST Guide. Also see `man 4 initrd`, and the documentation for the options `dhcp4 "create-cid"`, `dhcp6 "default-duid"` in `man 5 wicked-config`, `wicked duid --help`, and `wicked iaaid --help`.

18.2 Setting up a TFTP server

The following procedure describes how to prepare the server so that the client machines with UEFI and BIOS can boot remotely using files exported by TFTP.

18.2.1 Installing a TFTP server

To install a TFTP server, use the following procedure:

1. Install the `tftp` package.

```
> sudo zypper in tftp
```

2. Review the `tftpd` configuration in `/etc/sysconfig/tftp` and add or change options as required. Refer to `man 8 tftpd` for more details. The TFTP daemon works without changing the configuration. The default root directory for the files is `/srv/tftpboot`.
3. Ensure that `tftpd` is started at boot time, and restart it to read the new configuration.

```
> sudo systemctl enable tftp.socket
> sudo systemctl restart tftp.socket
```

18.2.2 Installing files for boot

SUSE Linux Enterprise Server provides the required files for booting via PXE on BIOS or UEFI machines. The following hardware architectures are supported:

- AMD64/Intel 64
- AArch64
- POWER
- IBM Z

Files required to boot from a specific hardware architecture are included in an RPM package. Install it on the machine running the TFTP server:

```
> sudo zypper in tftpboot-installation-SLE-OS_VERSION-ARCHITECTURE
```

Replace `OS_VERSION` with the version number of your SUSE Linux Enterprise Server installation, for example `SLE-15-SP3-x86_64`, and replace `ARCHITECTURE` with the architecture of your system, for example `x86_64`. So the resulting text would look like this: `tftpboot-installation-SLE-15-SP3-x86_64`. Run `zypper se tftpboot` to search for all available versions and architectures.

The files will be installed in `/srv/tftpboot/SLE-OS_VERSION-ARCHITECTURE`. You can also copy the files for other versions and architectures of SUSE Linux Enterprise Server to the `/srv/tftpboot` directory.



Tip: Serving different architectures

The client and server hardware architecture can vary. For example, you can run an AMD64/Intel 64 TFTP server and provide a bootable environment for AArch64 client machines by installing the `tftpboot-installation-SLE-15-SP3-aarch64` package.



Note: Existing `/srv/tftpboot/` directory

If the directory `/srv/tftpboot/` already exists on your machine, then all files will be installed to `/usr/share/tftpboot-installation/`. This is the case if you are upgrading your PXE server from a previous SLES release.

To fix this problem, copy the files manually from `/usr/share/tftpboot-installation/` to `/srv/tftpboot/`. Alternatively, remove `/srv/tftpboot/` and reinstall the `tftpboot-installation-SLE-OS_VERSION-ARCHITECTURE` package.

18.2.3 Configuring PXELINUX

Open the file `/srv/tftpboot/SLE-OS_VERSION-ARCHITECTURE/net/pxelinux.cfg/default` in an editor. Replace the path for the `install` parameter according to your setup as described in [Chapter 17, Setting up a network installation source](#). Also replace `TFTP_SERVER` with the IP address of the TFTP server. For an overview of the PXELINUX configuration options, see [Section 18.3, “PXELINUX configuration options”](#).

```
default linux

# install
label linux
    ipappend 2
    kernel boot/ARCHITECTURE/loader/linux
    append initrd=boot/ARCHITECTURE/loader/initrd instsys=tftp://TFTP_SERVER/
SLE-OS_VERSION-ARCHITECTURE/boot/ARCHITECTURE/root install=PROTOCOL://SERVER_IP:/PATH

display message
implicit 1
prompt 1
timeout 50
```

For details about the boot parameters that are used in the `append` line, see [Section 8.3, “List of important boot parameters”](#).

If required, edit the `/srv/tftpboot/SLE-OS_VERSION-ARCHITECTURE/net/pxelinux.cfg/message` to display a message in the boot menu.

18.2.4 Preparing PXE boot for EFI with GRUB2

Normally, the GRUB2 configuration files require no modifications. However, the default settings do not include a network resource for the installation system. To perform a full SUSE Linux Enterprise Server installation via network, you need to specify the `install` parameter in the `linuxefi` instruction of the `/srv/tftpboot/SLE-OS_VERSION-ARCHITECTURE/EFI/BOOT/grub.cfg` file. Refer to [Section 8.3.3, “Specifying the installation source”](#) for further information about the `install` parameter.

18.3 PXELINUX configuration options

The options listed here are a subset of all the options available for the PXELINUX configuration file.

APPEND OPTIONS

Adds one or more options to the kernel command line. These are added for both automatic and manual boots. The options are added at the very beginning of the kernel command line, usually permitting explicitly entered kernel options to override them.

APPEND -

Appends nothing. `APPEND` with a single hyphen as argument in a `LABEL` section can be used to override a global `APPEND`.

DEFAULT KERNEL_OPTIONS...

Sets the default kernel command line. When PXELINUX boots automatically, it executes the specified entries, appending the `auto` option.

If no configuration file exists or no `DEFAULT` entry is defined in the configuration file, the default is the kernel name “linux” with no options.

IFAPPEND FLAG

Adds a specific option to the kernel command line depending on the `FLAG` value. The `IFAPPEND` option is available only on PXELINUX. `FLAG` expects a value, described in [Table 18.1, “Generated and added kernel command line options from IFAPPEND”](#):

TABLE 18.1: GENERATED AND ADDED KERNEL COMMAND LINE OPTIONS FROM IFAPPEND

Argument	Generated kernel command line / Description
<u>1</u>	<code>ip=CLIENT_IP:BOOT_SERVER_IP:GW_IP:NETMASK</code>

Argument	Generated kernel command line / Description
	<p>The placeholders are replaced based on the input from the DHCP/BOOTP or PXE boot server.</p> <p>Note, this option is not a substitute for running a DHCP client in the booted system. Without regular renewals, the lease acquired by the PXE BIOS will expire, making the IP address available for reuse by the DHCP server.</p>
<u>2</u>	<p><code>BOOTIF=MAC_ADDRESS_OF_BOOT_INTERFACE</code></p> <p>This option is useful to avoid timeouts when the installation server probes one LAN interface after another until it gets a reply from a DHCP server. This option allows an initrd program to determine from which interface the system has been booted. linuxrc reads this option and uses this network interface.</p>
<u>4</u>	<p><code>SYSUUID=SYSTEM_UUID</code></p> <p>Adds UUIDs in lowercase hexadecimals, see /usr/share/doc/packages/syslinux/pxelinux.txt</p>

LABEL LABEL KERNEL IMAGE APPEND OPTIONS...

Indicates that if LABEL is entered as the kernel to boot, PXELINUX should instead boot IMAGE and the specified APPEND options should be used. They replace the ones specified in the global section of the file, before the first LABEL command. The default for IMAGE is the same as LABEL and, if no APPEND is given, the default is to use the global entry (if any). Up to 128 LABEL entries are permitted.

PXELINUX uses the following syntax:

```
label MYLABEL
    kernel MYKERNEL
    append MYOPTIONS
```

Labels are mangled as if they were file names and they must be unique after mangling. For example, the two labels “v2.6.30” and “v2.6.31” would not be distinguishable under PXELINUX because both mangle to the same DOS file name.

The kernel does not need to be a Linux kernel. It can also be a boot sector or a COMBOOT file.

LOCALBOOT *TYPE*

On PXELINUX, specifying `LOCALBOOT 0` instead of a `KERNEL` option means invoking this particular label and causes a local disk boot instead of a kernel boot.

Argument	Description
<u>0</u>	Perform a normal boot
<u>4</u>	Perform a local boot with the Universal Network Driver Interface (UNDI) driver still resident in memory
<u>5</u>	Perform a local boot with the entire PXE stack, including the UNDI driver, still resident in memory

All other values are undefined. If you do not know what the UNDI or PXE stacks are, specify `0`.

TIMEOUT *TIME-OUT*

Indicates how long to wait at the boot prompt until booting automatically, in units of 1/10 second. The time-out is canceled when the user types anything on the keyboard, assuming the user will complete the command begun. A time-out of zero disables the time-out completely (this is also the default). The maximum possible time-out value is 35996 (just less than one hour).

PROMPT *flag_val*

If `flag_val` is 0, displays the boot prompt only if `Shift` or `Alt` is pressed or `Caps Lock` or `Scroll Lock` is set (this is the default). If `flag_val` is 1, always displays the boot prompt.

```
F2  FILENAME
F1  FILENAME
..etc...
F9  FILENAME
F10 FILENAME
```

Displays the indicated file on the screen when a function key is pressed at the boot prompt. This can be used to implement preboot online help (presumably for the kernel command line options). For backward compatibility with earlier releases, `F10` can be also entered as `F0`. Note that there is currently no way to bind file names to `F11` and `F12`.

18.4 Preparing the target system for PXE boot

Prepare the system's BIOS for PXE boot by including the PXE option in the BIOS boot order.



Warning: BIOS boot order

Do not place the PXE option ahead of the hard disk boot parameter in the BIOS. Otherwise this system would try to re-install itself every time you boot it.

18.5 Using wake-on-LAN for remote wakeups

Wake-on-LAN (WOL) is an Ethernet standard for remotely waking up a computer by sending it a wakeup signal over a network. This signal is called the “magic packet”. Install WOL on client machines to enable remote wakeups, and on every machine you want to use for sending the wakeup signal. The magic packet is broadcast over UDP port 9 to the MAC address of the network interface on the client machine.

When computers are shut down they usually are not turned all the way off, but remain in a low power mode. When the network interface supports WOL, it listens for the magic packet wakeup signal when the machine is powered off. You can send the magic packet manually, or schedule wakeups in a cron job on the sending machine.

18.5.1 Prerequisites

WOL works with both wired and wireless Ethernet cards that support it.

You may need to enable WOL in your system BIOS/UEFI.

Check your BIOS/UEFI settings for PXE boot, and make sure it is disabled to prevent accidental re-installations.

Adjust your firewall to allow traffic over UDP port 9.

18.5.2 Verifying wired Ethernet support

Run the following command to see if a wired Ethernet interface supports WOL:

```
> sudo ethtool eth0 | grep -i wake-on
```

```
Supports Wake-on: pumbg
Wake-on: g
```

The example output shows that `eth0` supports WOL, indicated by the `g` flag on the `Supports Wake-on` line. `Wake-on: g` shows that WOL is already enabled, so this interface is ready to receive wakeup signals. If WOL is not enabled, enable it with this command:

```
> sudo ethtool -s eth0 wol g
```

18.5.3 Verifying wireless interface support

Wakeup-over-wifi, or WoWLAN, requires a wireless network interface that supports WoWLAN. Test it with the `iw` command, which is provided by the `iw` package:

```
> sudo zypper in iw
```

Find your device name:

```
> sudo iw dev
phy#0
    Interface wlan2
        ifindex 3
        wdev 0x1
        addr 9c:ef:d5:fe:01:7c
        ssid accesspoint
        type managed
        channel 11 (2462 MHz), width: 20 MHz, center1: 2462 MHz
        txpower 20.00 dBm
```

In this example, the device name to use for querying WoWLAN support is `phy#0`. This example shows that it is not supported:

```
> sudo iw phy#0 wowlan show
command failed: Operation not supported (-95)
```

This example shows an interface that supports WoWLAN, but is not enabled:

```
> sudo iw phy#0 wowlan show
WoWLAN is disabled
```

Enable it:

```
> sudo iw phy#0 wowlan enable magic-packet
WoWLAN is enabled:
```

```
* wake up on magic packet
```

18.5.4 Installing and testing WOL

To use WOL, install the `wol` package on the client and sending machines:

```
> sudo zypper in wol
```

Install `wol-udev-rules` on your client machines. This package installs a udev rule that enables WOL automatically at start-up.

Get the MAC address of the network interface on the client machine:

```
> sudo ip addr show eth0|grep ether  
link/ether 7c:ef:a5:fe:06:7c brd ff:ff:ff:ff:ff:ff
```

In the example output, `7c:ef:a5:fe:06:7c` is the MAC address.

Shut down your client machine, and send it a wakeup signal from another computer on the same subnet:

```
> wol 7c:ef:a5:fe:06:7c
```

If your target machine and second device are on the same network but in different subnets, specify the broadcast address for your target machine:

```
> wol -i 192.168.0.63 7c:ef:a5:fe:06:7c
```

Because WOL relies on broadcast domains, the sending machine must be on the same network, though it can be in a different network segment.

It is possible to send the magic packet from a different network. One way is with port forwarding, if your router supports port forwarding to a broadcast address. A more secure method is to connect to a host inside your network via SSH, and send the magic packet from there.

19 Setting up a UEFI HTTP Boot server

This chapter describes how to set up and configure a UEFI HTTP Boot server.

19.1 Introduction

HTTP Boot combines DHCP, DNS and HTTP to make it possible to boot and deploy systems over the network. HTTP Boot can be used as a high-performance replacement for PXE. HTTP Boot allows to boot a server from a URI over HTTP, quickly transferring large files, such as the Linux kernel and root file system, from servers outside of your local network.

19.1.1 Configuring the client machine

Enabling HTTP Boot on a physical client machine depends on your specific hardware. Consult the documentation for further information on how to enable HTTP Boot on your particular machine.

19.1.2 Preparation

The setup described here uses 192.168.111.0/24 (IPv4) and 2001:db8:f00f:cafe::/64 (IPv6) IP subnets and the server IP addresses are 192.168.111.1 (IPv4) and 2001:db8:f00f:cafe::1/64 (IPv6) as examples. Adjust these values to match your specific setup.

Install the following packages on the machine that you plan to use as an HTTP Boot server: `dhcp-server`, `apache2` (or `lighttpd`), and `dnsmasq`.

19.2 Configuring the server

19.2.1 DNS server

While configuring the DNS server is optional, this does allow you to assign a user-friendly name to the HTTP Boot server. To set up the DNS server, add the following to the `/etc/dnsmasq.conf` file:

```
interface=eth0
```

```
addn-hosts=/etc/dnsmasq.d/hosts.conf
```

Assign a domain name to the IP addresses in the `/etc/dnsmasq.d/hosts.conf` file:

```
192.168.111.1 www.httpboot.local
2001:db8:f00f:cafe::1 www.httpboot.local
```

Start the DNS server.

```
systemctl start dnsmasq
```



Note: Use the shim boot loader

Because of a change in UEFI 2.7, we recommend using a shim boot loader from SLE 15 or newer to avoid potential errors caused by the additional DNS node.

19.2.1.1 Configuring the DHCPv4 server

Before setting up the DHCP servers, specify the network interface for them in `/etc/sysconfig/dhcpd`:

```
DHCPD_INTERFACE="eth0"
DHCPD6_INTERFACE="eth0"
```

This way, the DHCP servers provide the service on the `eth0` interface only.

To set up a DHCPv4 server for both PXE Boot and HTTP Boot, add the following configuration to the `/etc/dhcpd.conf` file:

```
option domain-name-servers 192.168.111.1;
option routers 192.168.111.1;
default-lease-time 14400;
ddns-update-style none;
class "pxeclients" {
    match if substring (option vendor-class-identifier, 0, 9) = "PXEClient";
    option vendor-class-identifier "PXEClient";
    next-server 192.168.111.1;
    filename "/bootx64.efi";
}
class "httpclients" {
    match if substring (option vendor-class-identifier, 0, 10) = "HTTPClient";
    option vendor-class-identifier "HTTPClient";
    filename "http://www.httpboot.local/sle/EFI/BOOT/bootx64.efi";
}
```

```

}
subnet 192.168.111.0 netmask 255.255.255.0 {
    range dynamic-bootp 192.168.111.100 192.168.111.120;
    default-lease-time 14400;
    max-lease-time 172800;
}

```

Note that the DHCPv4 server must use the `HTTPClient` parameter for the vendor class ID, as the client uses it to identify an HTTP Boot offer.

Start the DHCP daemon:

```
systemctl start dhcpd
```

19.2.1.2 Configuring the DHCPv6 server

To set up the DHCPv6 server, add the following configuration to `/etc/dhcpd6.conf`:

```

option dhcp6.bootfile-url code 59 = string;
option dhcp6.vendor-class code 16 = {integer 32, integer 16, string};
subnet6 2001:db8:f00f:cafe::/64 {
    range6 2001:db8:f00f:cafe::42:10 2001:db8:f00f:cafe::42:99;
    option dhcp6.bootfile-url "http://www.httpboot.local/sle/EFI/B00T/bootx64.efi";
    option dhcp6.name-servers 2001:db8:f00f:cafe::1;
    option dhcp6.vendor-class 0 10 "HTTPClient";
}

```

This configuration defines the type of the boot URL, the vendor class, and other required options. Similar to the DHCPv4 settings, it is necessary to provide the boot URL, which must have an IPv6 address. It is also necessary to specify the vendor class option. In DHCPv6, it consists of the enterprise number and the vendor class data (length and the content). Since the HTTP Boot driver ignores the enterprise number, you can set it to `0`. The content of the vendor class data needs to be `HTTPClient`; otherwise, the client ignores the offer.

The older HTTP Boot implementation, which does not follow [RFC 3315](https://datatracker.ietf.org/doc/html/rfc3315), requires a different configuration:

```

option dhcp6.bootfile-url code 59 = string;
option dhcp6.vendor-class code 16 = string;
subnet6 2001:db8:f00f:cafe::/64 {
    range6 2001:db8:f00f:cafe::42:10 2001:db8:f00f:cafe::42:99;
    option dhcp6.bootfile-url "http://www.httpboot.local/sle/EFI/B00T/bootx64.efi";
    option dhcp6.name-servers 2001:db8:f00f:cafe::1;
    option dhcp6.vendor-class "HTTPClient";
}

```

Start the `dhcpv6` daemon.

```
systemctl start dhcpd6
```

19.2.1.2.1 Setting up the DHCPv6 server for both PXE and HTTP boot

Using the following configuration, it is possible to configure the DHCPv6 server for both PXE Boot and HTTP Boot:

```
option dhcp6.bootfile-url code 59 = string;
option dhcp6.vendor-class code 16 = {integer 32, integer 16, string};

subnet6 2001:db8:f00f:cafe::/64 {
    range6 2001:db8:f00f:cafe::42:10 2001:db8:f00f:cafe::42:99;

    class "PXEClient" {
        match substring (option dhcp6.vendor-class, 6, 9);
    }

    subclass "PXEClient" "PXEClient" {
        option dhcp6.bootfile-url "tftp://[2001:db8:f00f:cafe::1]/bootloader.efi";
    }

    class "HTTPClient" {
        match substring (option dhcp6.vendor-class, 6, 10);
    }

    subclass "HTTPClient" "HTTPClient" {
        option dhcp6.bootfile-url "http://www.httpboot.local/sle/EFI/B00T/bootx64.efi";
        option dhcp6.name-servers 2001:db8:f00f:cafe::1;
        option dhcp6.vendor-class 0 10 "HTTPClient";
    }
}
```

It is also possible to match the vendor class to a specific architecture, as follows:

```
class "HTTPClient" {
    match substring (option dhcp6.vendor-class, 6, 21);
}

subclass "HTTPClient" "HTTPClient:Arch:00016" {
    option dhcp6.bootfile-url "http://www.httpboot.local/sle/EFI/B00T/bootx64.efi";
    option dhcp6.name-servers 2001:db8:f00f:cafe::1;
    option dhcp6.vendor-class 0 10 "HTTPClient";
}
```

In this example, `HTTPClient:Arch:00016` refers to an AMD64/Intel 64 HTTP Boot client. This configuration allows the server to serve different architectures simultaneously.

19.2.1.2.2 Configuring firewall

If DHCPv6 packets are dropped by the RP filter in the firewall, check its log. In case it contains the `rpfilter_DROP` entry, disable the filter using the following configuration in `/etc/firewalld/firewalld.conf`:

```
IPv6_rpfilter=no
```

19.2.1.3 Deploying a TFTP server (optional)

To provide support for both PXE Boot and HTTP Boot, deploy a TFTP server. Install the `tftp` and start the service:

```
systemctl start tftp.socket
systemctl start tftp.service
```

It is also necessary to install a specific `tftpboot-installation` package for use with PXE Boot. Run the `zypper se tftpboot` command, to list of the available `tftp-installation` packages, then install the package for the desired system version and architecture, for example `tftpboot-installation-SLE-15-SP3-x86_64`. For example, `tftpboot-installation-SLE-VERSION-x86_64` (replace *VERSION* with the actual version). Copy the content of the `SLE-VERSION-x86_64` directory to the root directory of the TFTP server:

For more information, refer to `/usr/share/tftpboot-installation/SLE-VERSION-x86_64/README`

19.2.1.4 Setting up the HTTP server

Create the `sle/` directory under the `/srv/www/htdocs/` directory and copy the entire content of the first system ISO image to the `/srv/www/htdocs/sle/` directory. Then edit the `/srv/www/htdocs/sle/EFI/BOOT/grub.cfg` file. Use the following example as a reference:

```
timeout=60
default=1

menuentry 'Installation IPv4' --class opensuse --class gnu-linux --class gnu --class os {
```



```

    set gfxpayload=keep
    echo 'Loading kernel ...'
    linux /sle/boot/x86_64/loader/linux install=http://www.httpboot.local/sle
    echo 'Loading initial ramdisk ...'
    initrd /sle/boot/x86_64/loader/initrd
}

menuentry 'Installation IPv6' --class opensuse --class gnu-linux --class gnu --class os {
    set gfxpayload=keep
    echo 'Loading kernel ...'
    linux /sle/boot/x86_64/loader/linux install=install=http://www.httpboot.local/sle
    ipv6only=1 ifcfg=*=dhcp6,DHCLIENT6_MODE=managed
    echo 'Loading initial ramdisk ...'
    initrd /sle/boot/x86_64/loader/initrd
}

```

19.2.1.4.1 Configuring `lighttpd`

To enable the support for both IPv4 and IPv6 in `lighttpd`, modify `/etc/lighttpd/lighttpd.conf` as follows:

```

##
## Use IPv6?
##
#server.use-ipv6 = "enable"
$SERVER["socket"] == "[::]:80" { }

```

Start the `lighttpd` daemon:

```
systemctl start lighttpd
```

19.2.1.4.2 Configuring `apache2`

Apache requires no additional configuration. Start the `apache2` daemon:

```
systemctl start apache2
```

19.2.1.5 Enabling SSL support for the HTTP server (optional)

To use the HTTPS Boot, you need to convert an existing server certificate into the `DER` format and enroll it into the client's firmware.

Assuming you already have a certificate installed on your server, convert it into the DER format for use with the client using the following command:

```
openssl x509 -in CERTIFICATE.crt -outform der -out CERTIFICATE.der
```

19.2.1.5.1 Enroll the server certificate into the client firmware

The exact procedure of enrolling the converted certificate depends on the specific implementation of the client's firmware. For certain hardware, you need to enroll the certificate manually via the firmware UI using an external storage device with the certificate on it. Machines with Redfish support can enroll the certificate remotely. Consult the documentation for your specific hardware for more information on enrolling certificates.

19.2.1.5.2 Enabling SSL support in lighttpd

Since lighttpd needs the private key and the certificate in the same file, unify them using the following command:

```
cat CERTIFICATE.crt server.key > CERTIFICATE.pem
```

Copy CERTIFICATE.pem to the /etc/ssl/private/ directory.

```
cp server-almighty.pem /etc/ssl/private/  
chown -R root:lighttpd /etc/ssl/private/server-almighty.pem  
chmod 640 /etc/ssl/private/server-almighty.pem
```

Make sure that mod_openssl is listed in the server.modules section of the /etc/lighttpd/modules.conf file, for example:

```
server.modules = (  
    "mod_access",  
    "mod_openssl",  
)
```

Add the following lines to SSL Support section in /etc/lighttpd/lighttpd.conf:

```
# IPv4  
$SERVER["socket"] == ":443" {  
    ssl.engine          = "enable"  
    ssl.pemfile         = "/etc/ssl/private/server-almighty.pem"  
}
```

```
# IPv6
$SERVER["socket"] == "[::]:443" {
    ssl.engine          = "enable"
    ssl.pemfile         = "/etc/ssl/private/server-almighty.pem"
}
```

Restart `lighttpd` to activate SSL support:

```
systemctl restart lighttpd
```

19.2.1.5.3 Enabling SSL support in Apache

Open the `/etc/sysconfig/apache2` file and add the SSL flag as follows:

```
APACHE_SERVER_FLAGS="SSL"
```

Make sure that the `ssl` module is listed in `APACHE_MODULES`, for example:

Next, copy the private key and the certificate to the `/etc/apache2/` directory.

```
cp server.key /etc/apache2/ssl.key/
chown wwwrun /etc/apache2/ssl.key/server.key
chmod 600 /etc/apache2/ssl.key/server.key
cp server.crt /etc/apache2/ssl.crt/
```

Create the `ssl` vhost configuration.

```
cd /etc/apache2/vhosts.d
cp vhost-ssl.template vhost-ssl.conf
```

Edit `/etc/apache2/vhosts.d/vhost-ssl.conf` to change the private key and the certificate:

```
SSLCertificateFile /etc/apache2/ssl.crt/server.crt
SSLCertificateKeyFile /etc/apache2/ssl.key/server.key
```

Restart Apache to activate the SSL support:

```
systemctl restart apache2
```

19.2.1.5.4 Modify the DHCP configuration

Replace the `http://` prefix with `https://` in `dhcpd.conf/dhcpd6.conf` and restart the DHCP server.

```
systemctl restart dhcpd  
systemctl restart dhcpd6
```

19.3 Booting the client via HTTP boot

If the firmware already supports HTTP boot, plug in the cable and choose the correct boot option.

20 Deploying customized preinstallations

Rolling out customized preinstallations of SUSE Linux Enterprise Server to many identical machines spares you from installing each one of them separately and provides a standardized installation for the end users.

With YaST firstboot, create customized preinstallation images and determine the workflow for the final personalization steps that involve end user interaction (as opposed to AutoYaST, which allows completely automated installations).

Creating a custom installation, rolling it out to your hardware, and personalizing the final product involves the following steps:

1. Prepare the master machine whose disk needs to be cloned to the client machines. For more information, refer to [Section 20.1, “Preparing the master machine”](#).
2. Customize the firstboot workflow. For more information, refer to [Section 20.2, “Customizing the firstboot installation”](#).
3. Clone the master machine's disk and roll this image out to the clients' disks. For more information, refer to [Section 20.3, “Cloning the master installation”](#).
4. Have the end user personalize the instance of SUSE Linux Enterprise Server. For more information, refer to [Section 20.4, “Personalizing the installation”](#).

20.1 Preparing the master machine

To prepare a master machine for a firstboot workflow, proceed as follows:

1. Insert the installation media into the master machine.
2. Boot the machine.
3. Perform a normal installation including all necessary configuration steps, and make sure to select the `yast2-firstboot` package for installation.
4. To define your own workflow of YaST configuration steps for the end user or to add your own YaST modules to this workflow, proceed to [Section 20.2, “Customizing the firstboot installation”](#). Otherwise proceed directly to [Step 5](#).

5. Enable firstboot as root:

Create an empty file `/var/lib/YaST2/reconfig_system` to trigger firstboot's execution. This file will be deleted after the firstboot configuration has been successfully accomplished. Create this file using the following command:

```
touch /var/lib/YaST2/reconfig_system
```

6. Proceed to *Section 20.3, "Cloning the master installation"*.

20.2 Customizing the firstboot installation

Customizing the firstboot installation workflow may involve several components. Customizing them is recommended. If you do not make any changes, firstboot performs the installation using the default settings. The following options are available:

- Customizing messages to the user, as described in *Section 20.2.1, "Customizing YaST messages"*.
- Customizing licenses and license actions, as described in *Section 20.2.2, "Customizing the license action"*.
- Customizing the release notes to display, as described in *Section 20.2.3, "Customizing the release notes"*.
- Customizing the order and number of components involved in the installation, as described in *Section 20.2.4, "Customizing the workflow"*.
- Configuring additional optional scripts, as described in *Section 20.2.5, "Configuring additional scripts"*.

To customize any of these components, modify the following configuration files:

/etc/sysconfig/firstboot

Configure various aspects of firstboot (such as release notes, scripts, and license actions).

/etc/YaST2/firstboot.xml

Configure the installation workflow by enabling or disabling components or adding custom ones.

Provide translations for such a customized installation workflow, as described in *Section 20.2.6, "Providing translations of the installation workflow"*.



Tip: Alternative location of the control file

`/etc/YaST2/firstboot.xml` is the default path for the control file, installed by the `yast2-firstboot` package. If you need to define a different location for the control file, edit `/etc/sysconfig/firstboot`, and change the `FIRSTBOOT_CONTROL_FILE` variable to your preferred location.

If you want to customize more than the workflow components, refer to the `control.xml` documentation at https://doc.opensuse.org/projects/YaST/SLES11/tdg/inst_in_general_chap.html#product_control.

20.2.1 Customizing YaST messages

By default, an installation of SUSE Linux Enterprise Server contains several default messages that are localized and displayed at certain stages of the installation process. These include a welcome message, a license message, and a congratulatory message at the end of installation. You can replace any of these with your own versions and include localized versions of them in the installation. To include your own welcome message, proceed as follows:

1. Log in as `root`.
2. Open the `/etc/sysconfig/firstboot` configuration file and apply the following changes:
 - a. Set `FIRSTBOOT_WELCOME_DIR` to the directory path where you want to store the files containing the welcome message and the localized versions, for example:

```
FIRSTBOOT_WELCOME_DIR="/usr/share/firstboot/"
```

- b. If your welcome message has file names other than `welcome.txt` and `welcome_locale.txt` (where `locale` matches the ISO 639 language codes such as “cs” or “de”), specify the file name pattern in `FIRSTBOOT_WELCOME_PATTERNS`. For example:

```
FIRSTBOOT_WELCOME_PATTERNS="mywelcome.txt"
```

If unset, the default value of `welcome.txt` is assumed.

3. Create the welcome file and the localized versions and place them in the directory specified in the `/etc/sysconfig/firstboot` configuration file.

Proceed in a similar way to configure customized license and finish messages. These variables are `FIRSTBOOT_LICENSE_DIR` and `FIRSTBOOT_FINISH_FILE`.

Change the `SHOW_Y2CC_CHECKBOX` to “yes” if the user needs to be able to start YaST directly after performing the installation.

20.2.2 Customizing the license action

You can customize the way the installation system reacts to a user's refusal to accept the license agreement. There are three different ways in which the system could react to this scenario:

halt

The firstboot installation is aborted and the entire system shuts down. This is the default setting.

continue

The firstboot installation continues.

abort

The firstboot installation is aborted, but the system attempts to boot.

Make your choice and set `LICENSE_REFUSAL_ACTION` to the appropriate value.

20.2.3 Customizing the release notes

Depending on whether you have changed the instance of SUSE Linux Enterprise Server you are deploying with firstboot, you might need to educate the end users about important aspects of their new operating system. A standard installation uses release notes (displayed during one of the final stages of the installation) to provide important information to the users. To have your own modified release notes displayed as part of a firstboot installation, proceed as follows:

1. Create your own release notes file. Use the RTF format as in the example file in `/usr/share/doc/release-notes` and save the result as `RELEASE-NOTES.en.rtf` (for English).
2. Store optional localized versions next to the original version and replace the `en` part of the file name with the actual ISO 639 language code, such as `de` for German.
3. Open the firstboot configuration file from `/etc/sysconfig/firstboot` and set `FIRSTBOOT_RELEASE_NOTES_PATH` to the actual directory where the release notes files are stored.

20.2.4 Customizing the workflow

The provided `/etc/YaST2/firstboot.xml` example defines a standard workflow which includes the following enabled components:

- Language Selection
- Welcome
- License Agreement
- Time and Date
- Users
- Root Password
- Finish Setup

Bear in mind that this workflow is a template. You can adjust it properly by manually editing the firstboot configuration file `/etc/YaST2/firstboot.xml`. This XML file is a subset of the standard `control.xml` file that is used by YaST to control the installation workflow. See [Example 20.2, “Configuring the workflow section”](#) to learn more about how to configure the workflow section.

For an overview of proposals, see [Example 20.1, “Configuring the proposal screens”](#). This provides you with enough background to modify the firstboot installation workflow. The basic syntax of the firstboot configuration file (plus how the key elements are configured) is explained via this example.

EXAMPLE 20.1: CONFIGURING THE PROPOSAL SCREENS

```
...
<proposals config:type="list">❶
  <proposal>❷
    <name>firstboot_hardware</name>❸
    <mode>installation</mode>❹
    <stage>firstboot</stage>❺
    <label>Hardware Configuration</label>❻
    <proposal_modules config:type="list">❼
      <proposal_module>printer</proposal_module>❽
    </proposal_modules>
  </proposal>
</proposal>
...
</proposal>
</proposals>
```

- ❶ The container for all proposals that should be part of the firstboot workflow.
- ❷ The container for an individual proposal.
- ❸ The internal name of the proposal.
- ❹ The mode of this proposal. Do not make any changes here. For a firstboot installation, this must be set to installation.
- ❺ The stage of the installation process at which this proposal is invoked. Do not make any changes here. For a firstboot installation, this must be set to firstboot.
- ❻ The label to be displayed on the proposal.
- ❼ The container for all modules that are part of the proposal screen.
- ❽ One or more modules that are part of the proposal screen.

The next section of the firstboot configuration file consists of the workflow definition. All modules that should be part of the firstboot installation workflow must be listed here.

EXAMPLE 20.2: CONFIGURING THE WORKFLOW SECTION

```
<workflows config:type="list">
  <workflow>
    <defaults>
      <enable_back>yes</enable_back>
      <enable_next>yes</enable_next>
      <archs>all</archs>
    </defaults>
    <stage>firstboot</stage>
    <label>Configuration</label>
    <mode>installation</mode>
    ... <!-- list of modules -->
  </modules>
</workflow>
</workflows>
...
```

The overall structure of the workflows section is very similar to that of the proposals section. A container holds the workflow elements and the workflow elements all include stage, label and mode information (just as the proposals introduced in *Example 20.1, “Configuring the proposal screens”*). The most notable difference is the defaults section, which contains basic design information for the workflow components:

enable_back

Include the *Back* button in all dialogs.

enable_next

Include the *Next* button in all dialogs.

archs

Specify the hardware architectures on which this workflow should be used.

EXAMPLE 20.3: CONFIGURING THE LIST OF WORKFLOW COMPONENTS

```
<modules config:type="list">❶
  <module>❷
    <label>Language</label>❸
    <enabled config:type="boolean">>false</enabled>❹
    <name>firstboot_language</name>❺
  </module>
</modules>
```

- ❶ The container for all components of the workflow.
- ❷ The module definition.
- ❸ The label displayed with the module.
- ❹ The switch to enable or disable this component in the workflow.
- ❺ The module name. The module itself must be located under /usr/share/YaST2/clients.

To make changes to the number or order of proposal screens during the firstboot installation, proceed as follows:

1. Open the firstboot configuration file at /etc/YaST2/firstboot.xml.
2. Delete or add proposal screens or change the order of the existing ones:
 - To delete an entire proposal, remove the proposal element including all its sub-elements from the proposals section and remove the respective module element (with sub-elements) from the workflow.
 - To add a new proposal, create a new proposal element and fill in all the required sub-elements. Make sure that the proposal exists as a YaST module in /usr/share/YaST2/clients.
 - To change the order of proposals, move the respective module elements containing the proposal screens around in the workflow. Note that there may be dependencies on other installation steps that require a certain order of proposals and workflow components.
3. Apply your changes and close the configuration file.

You can always change the workflow of the configuration steps if the default does not meet your needs. Enable or disable certain modules in the workflow (or add your own custom ones).

To toggle the status of a module in the firstboot workflow, proceed as follows:

1. Open the `/etc/YaST2/firstboot.xml` configuration file.
2. Change the value for the `enabled` element from `true` to `false` to disable the module or from `false` to `true` to enable it again.

```
<module>
  <label>Time and Date</label>
  <enabled config:type="boolean">true</enabled>
  <name>firstboot_timezone</name>
</module>
```

3. Apply your changes and close the configuration file.

To add a custom made module to the workflow, proceed as follows:

1. Create your own YaST module and store the module file `module_name.rb` in `/usr/share/YaST2/clients`.
2. Open the `/etc/YaST2/firstboot.xml` configuration file.
3. Determine at which point in the workflow your new module should be run. In doing so, make sure that any dependencies on other steps in the workflow are taken into account and resolved.
4. Create a new `module` element inside the `modules` container and add the appropriate sub-elements:

```
<modules config:type="list">
  ...
  <module>
    <label>my_module</label>
    <enabled config:type="boolean">true</enabled>
    <name>filename_my_module</name>
  </module>
</modules>
```

- a. Enter the label to be displayed on your module in the `label` element.
- b. Make sure that `enabled` is set to `true` to have your module included in the workflow.

- c. Enter the file name of your module in the `name` element. Omit the full path and the `.rb` suffix.

5. Apply your settings and close the configuration file.



Tip: Finding connected network interface for auto-configuration

If the target hardware could feature more than one network interface, add the `network-autoconfig` package to the application image. `network-autoconfig` cycles through all available Ethernet interfaces until one is successfully configured via DHCP.

20.2.5 Configuring additional scripts

Firstboot can be configured to execute additional scripts after the firstboot workflow has been completed. To add additional scripts to the firstboot sequence, proceed as follows:

1. Open the `/etc/sysconfig/firstboot` configuration file and make sure that the path specified for `SCRIPT_DIR` is correct. The default value is `/usr/share/firstboot/scripts`.
2. Create your shell script, store it in the specified directory, and apply the appropriate file permissions.

20.2.6 Providing translations of the installation workflow

Depending on the end user it could be desirable to offer translations of the customized workflow. Those translations could be necessary if you customized the workflow by changing the `/etc/YaST2/firstboot.xml` file, as described in [Section 20.2.4, “Customizing the workflow”](#).

If you have changed `/etc/YaST2/firstboot.xml` and introduced string changes, generate a new translation template file (`.pot` file) and use the `gettext` toolchain to translate and finally install the translated files in the YaST locale directories (`/usr/share/YaST2/locale`) as compiled `.mo` files. Proceed as follows:

1. For example, change the `textdomain` setting from:

```
<textdomain>firstboot</textdomain>
```

to the following:

```
<textdomain>firstboot-oem</textdomain>
```

2. Use **xgettext** to extract the translatable strings to the translation template file (`.pot` file), for example to `firstboot-oem.pot`:

```
xgettext -L Glade -o firstboot-oem.pot /etc/YaST2/firstboot.xml
```

3. Start the translation process. Then package the translated files (`.LL_code.po` files) the same way as translations of the other projects and install the compiled `firstboot-oem.mo` files.

If you need translations for additional or changed YaST modules, provide translations within such a module itself. If you changed an existing module, make sure to change also its text-domain statement to avoid undesired side effects.



Tip: More information

For more information about YaST development, refer to https://en.opensuse.org/openSUSE:YaST_development⁷. Detailed information about YaST firstboot can be found at <https://doc.opensuse.org/projects/YaST/SLES11/tdg/bk09ch01s02.html>⁷.

20.3 Cloning the master installation

Clone the master machine's disk using any of the imaging mechanisms available to you, and roll these images out to the target machines. For more information about imaging, see <https://doc.suse.com/kiwi/>⁷.

20.4 Personalizing the installation



When the cloned disk image is booted, firstboot starts and the installation proceeds exactly as laid out in [Section 20.2.4, “Customizing the workflow”](#). Only the components included in the firstboot workflow configuration are started. All other installation steps are skipped. The end user adjusts language, keyboard, network, and password settings to personalize the workstation. After this process is finished, a firstboot installed system behaves as any other instance of SUSE Linux Enterprise Server.

A Imaging and creating products

To adapt the operating system better to your deployment, you can create custom media for use as an appliance or live system with KIWI NG. KIWI NG can be used either on a local machine or online in SUSE Studio Express (OBS).

With KIWI NG, you can create Live CDs, Live DVDs, flash disks to use on Linux-supported hardware platforms and virtual disks for virtualization and cloud systems (like Xen, KVM, VMware, EC2 and more). Images created by KIWI NG can also be used in a PXE environment to boot from the network.

This guide does not cover topics related to KIWI NG in depth, as there is separate documentation available:

- For more information, see the KIWI NG documentation at <https://doc.suse.com/kiwi/>  (also available in the package `kiwi-doc`).
- SUSE Studio Express on Open Build Service can be used to create OS images online. It supports creating virtual appliances and live systems, based on either openSUSE or SUSE Linux Enterprise. For more information and documentation, see <https://studioexpress.opensuse.org/> .

B 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. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

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

We have designed this License 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.