

SUSE Linux Enterprise Server

11 SP3

www.suse.com

Jun 03 2013

Deployment Guide



Deployment Guide

Copyright © 2006–2013 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 and Novell trademarks, see the Novell Trademark and Service Mark list <http://www.novell.com/company/legal/trademarks/tmlist.html>. All other third party trademarks are the property of their respective owners. A trademark symbol (®, TM etc.) denotes a SUSE or Novell trademark; an asterisk (*) denotes a third party trademark.

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

About This Guide xi

1 Available Documentation	xi
2 Feedback	xiii
3 Documentation Conventions	xiv

1 Planning for SUSE Linux Enterprise Server 1

1.1 Considerations for Deployment of a SUSE Linux Enterprise Server	2
1.2 Deployment of SUSE Linux Enterprise Server	3
1.3 Running SUSE Linux Enterprise Server	4

I Architecture Specific Installation Considerations 5

2 Installation on x86, AMD64, Intel 64, and Itanium 7

2.1 Required Background	7
2.2 System Requirements for Operating Linux	8
2.3 Installation Considerations	11
2.4 Boot and Installation Media	14
2.5 Installation Procedure	16
2.6 Controlling the Installation	17
2.7 Dealing with Boot and Installation Problems	19

3 Installation on IBM POWER	21
3.1 Requirements	21
3.2 Preparation	23
4 Installation on IBM System z	39
4.1 General Information and Requirements	39
4.2 Preparing for Installation	46
4.3 Network Connection Types	66
4.4 The parmfile—Automating the System Configuration	69
4.5 Using the vt220 Terminal Emulator	75
4.6 Further In-Depth Information about IBM System z	76
II Manual Deployment	79
5 Deployment Strategies	81
5.1 Deploying up to 10 Workstations	81
5.2 Deploying up to 100 Workstations	83
5.3 Deploying More than 100 Workstations	91
6 Installation with YaST	93
6.1 Choosing the Installation Method	93
6.2 The Installation Workflow	96
6.3 IBM POWER: System Start-Up for Network Installation	97
6.4 IBM System z: System Start-Up for Installation	97
6.5 System Start-Up for Installation	98
6.6 The Boot Screen on Machines Equipped with Traditional BIOS	98
6.7 The Boot Screen on Machines Equipped with UEFI	103
6.8 Welcome	105
6.9 IBM System z: Hard Disk Configuration	107

6.10 Media Check	109
6.11 Installation Mode	109
6.12 Clock and Time Zone	112
6.13 Server Base Scenario	114
6.14 Installation Settings	115
6.15 Performing the Installation	120
6.16 Configuration of the Installed System	123
6.17 Graphical Login	134

7 Updating SUSE Linux Enterprise 135

7.1 Terminology	135
7.2 The SUSE Linux Enterprise 11 Maintenance Model	137
7.3 Supported Upgrade Paths to SLE SP3	144
7.4 General Preparations for Updating	145
7.5 Updating SLE 11 SP1 to SLE 11 SP2	146
7.6 Updating SLE 11 SP2 to SLE 11 SP3	155
7.7 Backporting Source Code	160
7.8 The Atomic Update	163
7.9 Migration Hooks for YaST Wagon	165

8 Setting Up Hardware Components with YaST 171

8.1 Hardware Information	171
8.2 Setting Up Graphics Card and Monitor	172
8.3 Setting Up Keyboard and Mouse	173
8.4 Setting Up Sound Cards	175
8.5 Setting Up a Printer	179
8.6 Setting Up a Scanner	184

9 Installing or Removing Software	187
9.1 Definition of Terms	187
9.2 Using the KDE Interface (Qt)	189
9.3 Managing Software Repositories and Services	196
9.4 Keeping the System Up-to-date	199
10 Installing Add-On Products	205
10.1 Add-Ons	205
10.2 Binary Drivers	206
10.3 SUSE Software Development Kit (SDK) 11	206
11 Accessing the Internet	209
11.1 Direct Internet Connection	209
11.2 Internet Connection Via Network	211
12 Managing Users with YaST	213
12.1 User and Group Administration Dialog	213
12.2 Managing User Accounts	215
12.3 Additional Options for User Accounts	217
12.4 Changing Default Settings for Local Users	224
12.5 Assigning Users to Groups	224
12.6 Managing Groups	225
12.7 Changing the User Authentication Method	226
13 Changing Language and Country Settings with YaST	229
13.1 Changing the System Language	230
13.2 Changing the Country and Time Settings	234

14 Remote Installation 237

14.1 Installation Scenarios for Remote Installation	237
14.2 Setting Up the Server Holding the Installation Sources	246
14.3 Preparing the Boot of the Target System	255
14.4 Booting the Target System for Installation	266
14.5 Monitoring the Installation Process	270

15 Advanced Disk Setup 275

15.1 Using the YaST Partitioner	275
15.2 LVM Configuration	287
15.3 Soft RAID Configuration	293

16 Subscription Management 299

16.1 Using Kernel Parameters to Access an SMT Server	300
16.2 Configuring Clients Using AutoYaST Profile	301
16.3 Configuring Clients Using the clientSetup4SMT.sh Script	302
16.4 Registering Clients Against SMT Test Environment	303

III Imaging and Creating Products 305

17 KIWI 307

17.1 Prerequisites for KIWI	307
17.2 Knowing KIWI's Build Process	308
17.3 Image Description	308
17.4 Creating Appliances with KIWI	311
17.5 For More Information	313

18 Creating Add-On Products With Add-on Creator 315

18.1 Creating Images	315
18.2 Add-On Structure	316

18.3 For More Information	317
19 Creating Images with YaST Product Creator	319
19.1 Prerequisites for Product Creator	319
19.2 Creating Images	319
19.3 For More Information	321
20 Deploying Customized Preinstallations	323
20.1 Preparing the Master Machine	324
20.2 Customizing the Firstboot Installation	324
20.3 Cloning the Master Installation	333
20.4 Personalizing the Installation	333
IV Automated Installations	335
21 Automated Installation	337
21.1 Simple Mass Installation	337
21.2 Rule-Based Autoinstallation	349
21.3 For More Information	354
22 Automated Upgrade from SUSE Linux Enterprise 11 SP2 to 11 SP3	355
22.1 Preparing the AutoYaST Profile	355
22.2 Running the Automated Upgrade	357
22.3 GRUB Menu Section for Booting into the Upgrade	358
22.4 Second Stage of the Upgrade	358
22.5 Limitations and Hints	359
23 Automated Deployment of Preload Images	361
23.1 Deploying system manually from rescue image	362
23.2 Automated Deployment with PXE Boot	363

A GNU Licenses

369

A.1 GNU Free Documentation License	369
--	-----

About This Guide

Installations of SUSE Linux Enterprise Server are possible in many different ways. It is impossible to cover all combinations of boot, or installation server, automated installations or deploying images. This manual should help with selecting the appropriate method of deployment for your installation.

Part I, “Architecture Specific Installation Considerations” (page 5)

The standard deployment instructions differ depending on the architecture used. For differences and requirements regarding the architecture, see this part.

Part II, “Manual Deployment” (page 79)

Most tasks that are needed during installations are described here. This includes the manual setup of your computer as well as additional software and remote installations.

Part III, “Imaging and Creating Products” (page 305)

Mass installations often require the preparation of images or products furnished with the features that are needed in this special case. Several options are described that allow the administrator to prepare these deployment methods.

Part IV, “Automated Installations” (page 335)

To do unattended installations, either use the installation with AutoYaST or prepare an image with kiwi or firstboot. This part describes methods to deploy these installations with a minimum of user interaction.

Many chapters in this manual contain links to additional documentation resources, including additional documentation that is available on the system as well as documentation available on the Internet.

For an overview of the documentation available for your product and the latest documentation updates, refer to <http://www.suse.com/doc> or to the following section.

1 Available Documentation

We provide HTML and PDF versions of our books in different languages. The following manuals for users and administrators are available for this product:

Deployment Guide (page i)

Shows how to install single or multiple systems and how to exploit the product inherent capabilities for a deployment infrastructure. Choose from various approaches, ranging from a local installation or a network installation server to a mass deployment using a remote-controlled, highly-customized, and automated installation technique.

Administration Guide (↑*Administration Guide*)

Covers system administration tasks like maintaining, monitoring, and customizing an initially installed system.

Security Guide (↑*Security Guide*)

Introduces basic concepts of system security, covering both local and network security aspects. Shows how to make use of the product inherent security software like AppArmor (which lets you specify per program which files the program may read, write, and execute), and the auditing system that reliably collects information about any security-relevant events.

Security and Hardening (↑*Security and Hardening*)

Deals with the particulars of installing and setting up a secure SUSE Linux Enterprise Server, and additional post-installation processes required to further secure and harden that installation. Supports the administrator with security-related choices and decisions.

System Analysis and Tuning Guide (↑*System Analysis and Tuning Guide*)

An administrator's guide for problem detection, resolution and optimization. Find how to inspect and optimize your system by means of monitoring tools and how to efficiently manage resources. Also contains an overview of common problems and solutions, and of additional help and documentation resources.

Virtualization with Xen (↑*Virtualization with Xen*)

Offers an introduction to virtualization technology of your product. It features an overview of the various fields of application and installation types of each of the platforms supported by SUSE Linux Enterprise Server as well as a short description of the installation procedure.

Virtualization with KVM for IBM System z (↑*Virtualization with KVM for IBM System z*)

Offers an introduction to setting up and managing virtualization with KVM (Kernel-based Virtual Machine) on SUSE Linux Enterprise Server. Learn

how to manage KVM with libvirt or QEMU. The guide also contains detailed information about requirements, limitations, and support status.

AutoYaST (↑AutoYaST)

AutoYaST is a system for installing one or more SUSE Linux Enterprise systems automatically and without user intervention, using an AutoYaST profile that contains installation and configuration data. The manual guides you through the basic steps of auto-installation: preparation, installation, and configuration.

Storage Administration Guide (↑Storage Administration Guide)

Provides information about how to manage storage devices on a SUSE Linux Enterprise Server.

In addition to the comprehensive manuals, several quick start guides are available:

Installation Quick Start (↑Installation Quick Start)

Lists the system requirements and guides you step-by-step through the installation of SUSE Linux Enterprise Server from DVD, or from an ISO image.

Linux Audit Quick Start

Gives a short overview how to enable and configure the auditing system and how to execute key tasks such as setting up audit rules, generating reports, and analyzing the log files.

AppArmor Quick Start

Helps you understand the main concepts behind AppArmor®.

Virtualization with Linux Containers (LXC) (↑Virtualization with Linux Containers (LXC))

Gives a short introduction to LXC (a lightweight “virtualization” method) and shows how to set up an LXC host and LXC containers.

Find HTML versions of most product manuals in your installed system under `/usr/share/doc/manual` or in the help centers of your desktop. Find the latest documentation updates at <http://www.suse.com/doc> where you can download PDF or HTML versions of the manuals for your product.

2 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests

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

To report bugs for a product component, log in to the Novell Customer Center from <http://www.suse.com/support/> and select *My Support > Service Request*.

User Comments

We want to hear your comments about and suggestions for this manual and the other documentation included with this product. Use the User Comments feature at the bottom of each page in the online documentation or go to <http://www.suse.com/doc/feedback.html> and enter your comments there.

Mail

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

3 Documentation Conventions

The following typographical conventions are used in this manual:

- `/etc/passwd`: directory names and filenames
- *placeholder*: replace *placeholder* with the actual value
- `PATH`: the environment variable `PATH`
- `ls, --help`: commands, options, and parameters
- `user`: users or groups
- `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

- **amd64 em64t ipf:** This paragraph is only relevant for the architectures `amd64`, `em64t`, and `ipf`. The arrows mark the beginning and the end of the text block.
- **ipseries zseries:** This paragraph is only relevant for the architectures `System z` and `ipseries`. The arrows mark the beginning and the end of the text block.
- *Dancing Penguins* (Chapter *Penguins*, ↑Another Manual): This is a reference to a chapter in another manual.

Planning for 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. SUSE Linux Enterprise Server 11 SP3 provides a variety of new features. It is impossible to describe all the new features here. The following is just a list of major enhancements that might be of interest.

Xen 4.0 Virtualization

Runs many virtual machines on a single server, each with its own instance of an operating system. For more information, see *Virtualization with Xen* (↑*Virtualization with Xen*).

YaST

Several new configuration options have been developed for YaST. These are normally described in the chapters about the technology involved.

SPident

The management utility SPident gives an overview of the installed software base and clarifies the current service pack level of the system.

Directory Services

Several LDAP-compliant directory services are available:

- Microsoft Active Directory
- OpenLDAP

AppArmor

Harden your System with the AppArmor technology. This service is described in depth in Part “Confining Privileges with AppArmor” (↑*Security Guide*).

AIDE

This is an intrusion detection system that can be setup to detect unauthorized changes to the system.

iSCSI

iSCSI provides an easy and reasonably inexpensive solution for connecting Linux computers to central storage systems. Find more information about iSCSI in Storage Administration Guide (↑*Storage Administration Guide*).

Network File System v4

Starting with version 10, SUSE Linux Enterprise Server supports NFS also in version 4. This gives you performance improvements, strong security, and a “stateful” protocol.

Oracle Cluster File System 2

OCFS2 is a general-purpose journaling file system that is fully integrated in the Linux 2.6 kernel and later. Find an overview of OCFS2 in the *High Availability Guide*.

Linux Kernel Crash Dump

Debugging kernel-related problems is now much easier when using Kexec and Kdump. This technology is available on x86, AMD64, Intel 64, and POWER platforms.

1.1 Considerations for Deployment of a SUSE Linux Enterprise Server

At the beginning of the planning process, you should try to define the project goals and needed 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. See also Chapter 5, *Deployment Strategies* (page 81).
- Will the system run as physical host or as a virtual machine?

- Will the system be in a hostile environment? Have a look at Chapter 1, *Security and Confidentiality* (↑*Security Guide*) to get an overview of consequences.
- How will you get regular updates? All patches are provided online for registered users. Find the registration and patch support database at <http://download.novell.com/patch/finder/>.
- Do you need help for your local installation? Novell provides training, support, and consulting for all topics pertaining to SUSE Linux Enterprise Server. Find more information about this at <http://www.novell.com/products/server/>.
- Do you need third-party products? Make sure that the required product is also supported on the desired platform. Novell 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 <http://developer.novell.com/yessearch/Search.jsp>.

Depending on the number of desired installations, it is beneficial to use installation servers or even completely automatic installations. Have a look at Chapter 5, *Deployment Strategies* (page 81) for more information. When using Xen 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. Find an overview of the documentation in this book in About This Guide (↑*Administration Guide*). 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 as well as help desk staff.

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. For any serious computing task where data loss could occur, a regular backup should be done.

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 where you can apply all changes for testing purposes before doing so in production. This also gives you the possibility of switching machines in the case of hardware failure.

Part I. Architecture Specific Installation Considerations

Installation on x86, AMD64, Intel 64, and Itanium

This chapter describes the steps necessary to prepare for the installation of SUSE Linux Enterprise Server on x86, AMD64, Intel 64, and Itanium computers. It introduces the steps required to prepare for various installation methods. The list of hardware requirements provides an overview of supported 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.

2.1 Required Background

To keep the scope of these guidelines manageable, certain technical assumptions have been made:

- You have some computer experience and are familiar with common technical terms.
- You are familiar with the documentation for your system and the network on which it runs.
- You have a basic understanding of Linux systems.

For an overview of the documentation available for your product and the latest documentation updates, refer to <http://www.suse.com/doc>.

2.2 System Requirements for Operating Linux

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 <http://developer.novell.com/yessearch/Search.jsp>.

2.2.1 Hardware for x86

Computers based on x86 constitute a cost-effective way of building high-performance systems. The preconditions for operating SUSE Linux Enterprise Server on this platform are as follows:

CPU

The number of CPUs supported depends on the kernel used. Specifically, these are as follows:

Table 2.1: *CPUs Supported by the Kernel*

Kernel	Oldest CPU Type	Maximum Number of CPUs
kernel-default	PentiumPro, Athlon	32
kernel-pae	Pentium II, Athlon XP	128

Memory Requirements

A minimum of 512 MB is required. The recommended memory is 1 GB. For a multiprocessor system, 256 MB per processor is required. Systems with less than 1 GB main memory need additional swap space that extends the virtual memory to 1 GB.

Hard Disk Requirements

The disk requirements depend largely on the installation. Commonly, you need more space than the installation software itself needs to have a system that works properly. Minimal requirements for different selections are:

System	Hard Disk Requirements
Minimal X Window System	1.2 GB
GNOME Desktop	3.2 GB
KDE Desktop	2.7 GB
All patterns	10 GB

Boot Methods

The computer can be booted for installation from DVD, USB hard drive, or the network. A special boot server is required to boot over the network. This boot server can be configured with SUSE Linux Enterprise Server. To use USB hard drives, the BIOS or firmware must support booting from USB devices. Create a bootable USB hard drive as described in Table 6.1, “Boot Options” (page 94).

2.2.2 Hardware for Itanium

The Itanium architecture is 64-bit and allows the operation of large servers.

CPU

II (older Itanium CPUs are no longer supported). Dual core CPUs and hyperthreading are also supported.

Maximum Number of CPUs

At most, 4096 CPUs are supported. For the calculation of the CPU count, a dual-core CPU counts as two CPUs and a hyperthreaded CPU with two siblings also counts as two CPUs. 1024 CPUs could mean 512 dual cores, 512 single cores with hyperthreading, or 256 dual cores with hyperthreading.

Memory

A minimum of 1GB RAM per CPU socket is recommended.

Hard Disk Requirements

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

System	Hard Disk Requirements
Minimal System	4 GB
Recommended	10 GB

Boot Methods

Options for booting the computer depend on the available hardware. All boot methods available to the machine should work. A special boot server is required to use PXE boot over the network. This may also be set up with SUSE Linux Enterprise Server.

2.2.3 Hardware for AMD64 and Intel 64

The AMD64 and Intel 64 architectures support the simple migration of x86 software to 64 bits. Like the x86 architecture, they constitute a value-for-money alternative.

CPU

All CPUs available on the market to date are supported. This includes dual-core CPUs.

Maximum Number of CPUs

The maximum number of CPUs supported by AMD64 and Intel 64 is 128.

Memory Requirements

A minimum of 512 MB of memory is required. Requirements depend on the application. However, the minimum recommended is 1024 MB or 512 MB per CPU on multiprocessor computers. The theoretical upper limit on the amount of memory supported by the kernel is 512 GB.

Hard Disk Requirements

The disk requirements depend largely on the installation selected. The required space for this architecture is similar to x86 but you should allocate some space for compatibility libraries. Minimum requirements for different selections are:

System	Hard Disk Requirements
Minimal X Window System	1.4 GB
GNOME Desktop	3.5 GB
KDE Desktop	3 GB
All patterns	8.5 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.

2.2.4 Supported Virtualization Hosts

The i586 and x86_64 version of SUSE Linux Enterprise Server can also be installed as VM Guests on various virtualization hosts. The following host operating systems and virtualization platforms are supported:

- KVM on SLES 11 SP2+
- XEN on SLES 10 SP4 / 11 SP1+
- Citrix XenServer 6.0 / 6.1
- Microsoft Windows 2008 SP2+ / 2008 R2+ / 2012+
- Oracle VM 3.0 / 3.1 / 3.2
- VMware ESX 5.1 / ESXi 5.1 / ESX 5.2 / ESXi 5.2

2.3 Installation Considerations

This section encompasses many factors that need to be considered before installing SUSE Linux Enterprise Server on x86, AMD64, Intel 64, and Itanium hardware.

2.3.1 Installation Type

SUSE Linux Enterprise Server is normally installed as an independent operating system. With the introduction of Xen, it is also possible to run multiple instances of SUSE Linux Enterprise Server on the same hardware. However, the controlling Domain-0 installation for Xen is performed like a typical installation with some additional packages. The installation of Xen guests is described in Chapter 3, *Setting Up Virtual Machines* (↑*Virtualization with Xen*).

2.3.2 Boot Methods

Depending on the hardware used, the following boot methods are available for the first boot procedure (prior to the installation of SUSE Linux Enterprise Server).

Table 2.2: *Boot Options*

Boot Option	Use
CD or DVD drive	The simplest booting method. The system requires a locally-available CD-ROM or DVD-ROM drive for this.
Floppy or USB disks	Find the images required for creating boot disks on the first CD or DVD in the <code>/boot</code> directory. See also the <code>README</code> in the same directory. Booting from a USB memory stick is only possible if the BIOS of the machine supports this method.
PXE or bootp	Must be supported by the BIOS or by the firmware of the system used. This option requires a boot server in the network. This task can be handled by a separate SUSE Linux Enterprise Server.

Boot Option	Use
Hard disk	SUSE Linux Enterprise Server can also be booted from hard disk. For this, copy the kernel (<code>linux</code>) and the installation system (<code>initrd</code>) from the <code>/boot/loader</code> directory of the first CD or DVD onto the hard disk and add an appropriate entry to the boot loader.

2.3.3 Installation Source

When installing SUSE Linux Enterprise Server, the actual installation data must be available in the network, on a hard disk partition, or on a local DVD. To install from the network, you need an installation server. To make the installation data available, set up any computer in a Unix or Linux environment as an NFS, HTTP, SMB, or FTP server. To make the installation data available from a Windows computer, release the data with SMB.

The installation source is particularly easy to select if you configure an *SLP server* in the local network. For more information, see Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246).

2.3.4 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 are useful if there is a need to start the

same system in different locations or if you intend to use Xen features like domain migration.

2.3.5 Different Installation Methods

SUSE Linux Enterprise Server offers several different methods for controlling installation:

- Installation on the console
- Installation via serial console
- Installation with AutoYaST
- Installation with KIWI images
- Installation via SSH
- Installation with VNC

By default, the graphical console is used. If you have a large number of similar computers to install, it is advisable to create an AutoYaST configuration file or a KIWI preload image and make this available to the installation process. See also the documentation for `autoyast2` in Chapter 21, *Automated Installation* (page 337) and `KIWI` in Chapter 17, *KIWI* (page 307).

2.4 Boot and Installation Media

When installing the system, the media for booting and for installing the system may be different. All combinations of supported media for booting and installing may be used.

2.4.1 Boot Media

Booting a computer depends on the capabilities of the hardware used and the availability of media for the respective boot option.

Booting from DVD

This is the most common possibility of booting a system. It is straightforward for most computer users, but requires a lot of interaction for every installation process.

Booting from a USB Hard Drive

Depending on the hardware used, it is possible to boot from a USB hard drive. The respective media must be created as described in Table 6.1, “Boot Options” (page 94).

Booting from the Network

You can only boot a computer directly from the network if this is supported by the computer's firmware or BIOS. 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 pxeboot. If you need a boot server, also read Section 14.1.3, “Remote Installation via VNC—PXE Boot and Wake on LAN” (page 241).

2.4.2 Installation Media

The installation media contain all the necessary packages and meta information that is necessary to install a SUSE Linux Enterprise Server. These must be available to the installation system after booting for installation. Several possibilities for providing the installation media to the system are available with SUSE Linux Enterprise Server.

Installation from DVD

All necessary data is delivered on the boot media. Depending on the selected installation, a network connection or add on media may be necessary.

Networked Installation

If you plan to install several systems, providing the installation media over the network makes things a lot easier. It is possible to install from many common protocols, such as NFS, HTTP, FTP, or SMB. For more information on how to run such an installation, refer to Chapter 14, *Remote Installation* (page 237).

2.5 Installation Procedure

This section offers an overview of the steps required for the complete installation of SUSE® Linux Enterprise Server in the required mode. Part II, “Manual Deployment” (page 79) contains a full description of how to install and configure the system with YaST.

2.5.1 Booting from a Local Interchangeable Drive

CD-ROM and floppy drives and USB memory sticks can be used for installation purposes. Adjust your computer to your needs:

1. Make sure that the drive is entered as a bootable drive in the BIOS.
2. Insert the boot medium in the drive and start the boot procedure.
3. The boot menu of the CD, DVD, floppy, or USB disk allows transferring different parameters to the installation system. See also Section 14.4.2, “Using Custom Boot Options” (page 267). If the installation should be performed over the network, specify the installation source here.
4. If unexpected problems arise during installation, use safe settings to boot.

2.5.2 Installing over the Network

An installation server is required to perform the installation by using a network source. The procedure for installing this server is outlined in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246).

If you have an SLP server, select SLP as the installation source in the first boot screen. During the boot procedure, select which of the available installation sources to use.

If the DVD is available on the network, use it as an installation source. In this case, specify the parameter `install=<URL>` with suitable values at the boot prompt. Find a more detailed description of this parameter in Section 14.4.2, “Using Custom Boot Options” (page 267).

2.6 Controlling the Installation

Control the installation in one of several ways. The method most frequently used is to install SUSE® Linux Enterprise Server from the computer console. Other options are available for different situations. Find more information about the available installation methods in Chapter 5, *Deployment Strategies* (page 81).

2.6.1 Installation on the Computer Console

The simplest way to install SUSE Linux Enterprise Server is using the computer console. With this method, a graphical installation program guides you through the installation. This installation method is discussed in detail in Chapter 6, *Installation with YaST* (page 93).

You can still perform the installation on the console without a working graphics mode. The text-based installation program offers the same functionality as the graphical version. Find some hints about navigation in this mode in Section “Navigation in Modules” (Chapter 3, *YaST in Text Mode*, ↑*Administration Guide*).

2.6.2 Installation Using a Serial Console

For this installation method, you need a second computer that is connected by a *null modem* cable to the computer on which to install SUSE Linux Enterprise Server. Depending on the hardware, even the firmware or BIOS of the computer may already be accessible to the serial console. If this is possible, you can carry out the entire installation using this method. To activate the serial console installation, additionally specify the parameter `console=ttyS0` at the boot prompt after the boot process has completed and before the installation system starts.

On most computers, there are two serial interfaces, *ttyS0* and *ttyS1*. 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/ttyS0 9600
```

This means that screen listens to the first serial port with a baud rate of 9600. From this point on, the installation proceeds similarly to the text-based installation over this terminal.

2.6.3 Installation with SSH

If you do not have direct access to the computer hardware and, for example, the installation should be launched from a management console, control the entire installation process over the network. To do this, enter the parameters `UseSSH=1` and `SSHPassword=<secret>` at the boot prompt. An SSH daemon is then launched in the system and you can log in to the system as user `root` with the password “secret”. To connect, use the command `ssh -X root@<ipaddr>`.

If you do not have a DHCP server available in your local network, manually assign an IP address to the installation system. Do this by entering the option `HostIP=<ipaddr>` at the boot prompt.

As soon as you are logged in to the installation system, launch the actual installation with the command `yast`. The installation will start in the graphical mode if `DISPLAY` is set. This then guides you through the installation. This procedure is described in detail in Section 14.1.5, “Simple Remote Installation via SSH—Dynamic Network Configuration” (page 243).

2.6.4 Installation over VNC

If you do not have direct access to the system, but want a graphical installation, install SUSE Linux Enterprise Server over VNC. This method is described in detail in Section 14.5.1, “VNC Installation” (page 270).

As suitable VNC clients are also available for other operating systems, such as Microsoft Windows and MacOS, the installation can also be controlled from computers running those operating systems.

2.6.5 Installation with AutoYaST

If you need to install SUSE Linux Enterprise Server on a number of computers with similar hardware, it is recommended you perform the installations with the aid of

AutoYaST. In this case, start by installing one SUSE Linux Enterprise Server and use this to create the necessary AutoYaST configuration files.

AutoYaST is extensively documented in Chapter 21, *Automated Installation* (page 337).

2.7 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.

2.7.1 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.

Installed System Boots, Not 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 are unable to 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 2.6, “Controlling the Installation” (page 17).

The Itanium Boot Loader

If you have manually altered the kernel or initrd on your system, run `/sbin/elilo` before shutting down the computer. If you leave out this step, your system may not be bootable.

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. A boot disk may also be a suitable solution in connection with OS/2 or Windows NT. For more details on creating boot disks, see Section “Creating Boot CDs” (Chapter 10, *The Boot Loader GRUB*, ↑*Administration Guide*).

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 or LILO. 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.

2.7.2 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 in failsafe mode. 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`.

2.7.3 Redirecting the Boot Source to the Boot DVD

To facilitate the installation process and avoid accidental installations, the default setting on the installation DVD for SUSE Linux Enterprise Server is that your system is booted from the first hard disk. At this point, an installed boot loader normally takes over control of the system. This means that the boot DVD can stay in the drive during an installation. To start the installation, choose one of the installation possibilities in the boot menu of the media.

Installation on IBM POWER

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

3.1 Requirements

A standard installation requires at least 256 MB of RAM. The installation of a standard system requires at least 2.0 GB of free hard disk space.

3.1.1 Hardware Requirements

The SUSE® Linux Enterprise Server operating system can be operated on a wide range of hardware. 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, check the database of hardware certified by SUSE. Find a list of certified hardware at <http://developer.novell.com/yessearch/Search.jsp>.

SUSE Linux Enterprise Server may support additional IBM POWER systems not listed below. For the latest information, see the IBM Information Center for Linux at <http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaam%2Fliaamdistros.htm>.

Find up-to-date firmware at IBM FixCentral (<http://www.ibm.com/support/fixcentral/>). Select your system from the Product Group list.

All systems listed below are operated with a PPC64 kernel.

3.1.1.1 IBM POWER7 Systems

POWER7 systems

- Power 710 Express
- Power 720 Express
- Power 730 Express
- Power 740 Express
- Power 750 Express
- Power 755
- Power 770
- Power 780
- Power 795

POWER7 BladeCenter Models

- IBM BladeCenter PS700
- IBM BladeCenter PS701
- IBM BladeCenter PS702
- IBM BladeCenter PS703
- IBM BladeCenter PS704

3.1.1.2 IBM PowerLinux Systems

- IBM PowerLinux 7R2

3.1.1.3 IBM POWER5 and POWER6 Systems

POWER5 Systems

- OpenPower 710

- System p5 520
- System i5 520

POWER6 Systems

- IBM Power 550
- IBM Power 570

POWER6 BladeCenter models

- IBM BladeCenter JS12
- IBM BladeCenter JS22
- IBM BladeCenter JS23
- IBM BladeCenter JS43

3.2 Preparation

This section describes the preparatory steps that must be taken before the actual installation. The installation procedure depends on the system used. See the following documentation:

- For IBM eServer p5 Systems, see Section 3.2.1, “Preparing for Installation on IBM eServer p5, System p, and OpenPower Models” (page 24)
- For IBM pSeries systems, see Section 3.2.2, “Preparing for Installation on IBM pSeries Models” (page 31)
- For IBM JS20/JS21/JS22 Blades, see Section 3.2.3, “Preparing an Installation on IBM JSxx BladeCenter” (page 35)

If SUSE® Linux Enterprise Server needs to be installed on a number of systems or partitions, it is recommended you create a network installation source. The same source can also be used for the concurrent installation on several partitions or several systems. The configuration of a network installation source is described in Section 14.2.1, “Setting Up an Installation Server Using YaST” (page 246).

The installation can be controlled with a VNC client. For more information about VNC, see Section 14.1.1, “Simple Remote Installation via VNC—Static Network Configuration” (page 238).

To participate in the `linuxppc-dev` mailing list, sign up using the forms at <http://lists.ozlabs.org/listinfo/linuxppc-dev/>. The following links are pertinent to the maintenance of an installation:

- <http://www.novell.com/support/products/server/> is an effective help tool for assisting customers in solving problems. A corresponding article is published whenever SUSE discover that a special case could lead to serious problems. Search the portal using keywords like PPC or POWER.
- Find security alerts at <http://www.suse.com/support/security/>. SUSE also maintains two security-related mailing lists to which anyone may subscribe.
 - `suse-security` — General discussion of security regarding 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.

3.2.1 Preparing for Installation on IBM eServer p5, System p, and OpenPower Models

This section covers the preparatory steps for installing SUSE® Linux Enterprise Server on IBM eServer p5 systems. It explains the installation from a built-in CD-ROM drive and over the network.

This section assumes you have set up your HMC and connected it to your system. Find more information about using the wizard to configure the HMC in “Configuring the HMC using the Guided Setup Wizard”: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphai_p5/confighmcgs.htm?

3.2.1.1 Modern Features of IBM eServer p5 Systems

IBM eServer p5 systems offer the possibility of partitioning the system. This enables the concurrent operation of up to 254 operating systems on one machine. These operating systems are installed in *LPARs* (logical partitions). One or several of these partitions can contain a SUSE Linux Enterprise Server environment.

To prepare an LPAR for SUSE Linux Enterprise Server, first configure the system over the *HMC*. Refer to the IBM documentation for details: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphbi/iphbikickoff.htm>


3.2.1.2 Hard Disk Space

Make sure that you have sufficient hard disk space for installing SUSE Linux Enterprise Server. The standard system requires at least 4 GB of free hard disk space.

3.2.1.3 Assigning an Installation Device to an LPAR

SUSE Linux Enterprise Server can be installed from a CD-ROM or DVD drive or using a network installation source. Make the CD-ROM, DVD drive, or network device available to the LPAR to install.

Figure 3.1: *HMC: Server Management—Properties*

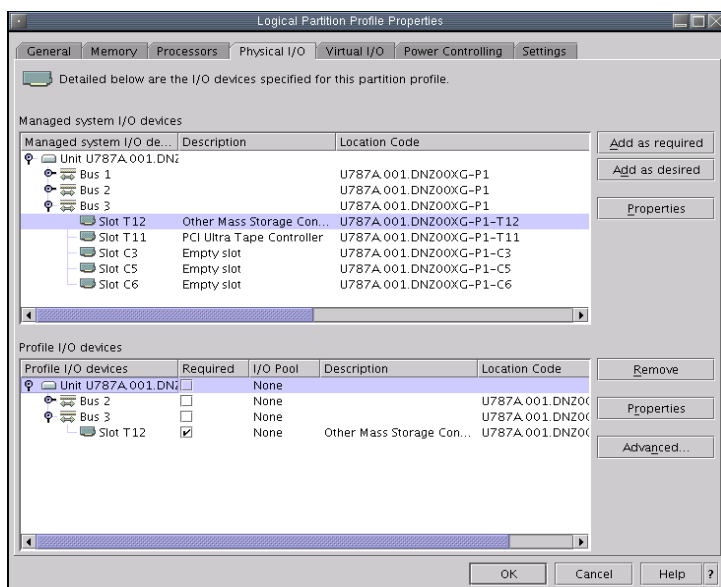


Procedure 3.1: *Assigning a CD-ROM or DVD Drive to an LPAR*

- 1 Open the HMC application and go to *Server and Partition > Server Management*.
- 2 From the available servers, expand the server and partition to install.
- 3 Right-click the profile to use for installation and select *Properties*—see Figure 3.1, “HMC: Server Management—Properties” (page 26).
- 4 In the *Logical Partition Profile Properties* dialog, select the *Physical I/O* tab.
- 5 From *Managed system I/O devices*, select the *Other Mass Storage Controller* from the bus where it is installed. To assign this DVD drive to the partition, click *Add as required*.

The result should look like Figure 3.2, “HMC: Managed System I/O Devices” (page 27).

Figure 3.2: *HMC: Managed System I/O Devices*



Now insert the SUSE Linux Enterprise Server CD1 or DVD1 in the drive.

Procedure 3.2: *Assigning a Network Device to an LPAR*

- 1 Open the HMC application and go to *Server and Partition > Server Management*.
- 2 From the available servers, open the server and partition to install.
- 3 Right-click the profile to use for installation and select *Properties*—see Figure 3.1, “HMC: Server Management—Properties” (page 26).
- 4 In the *Logical Partition Profile Properties* dialog, select the *Physical I/O* tab.
- 5 From *Managed system I/O devices*, select *PCI 10/100/1000Mbps Ethernet UTP 2-port* from the bus where it is installed. Then click *Add as required*.

If you plan to install using a virtual ethernet adapter, refer to the IBM documentation.

Create a network installation source if SUSE Linux Enterprise Server should be installed on a number of partitions. This eliminates the need to change CDs

during installation. The same source can also be used for concurrent installation of various systems. The configuration of the network installation source is described in Section 14.2.1, “Setting Up an Installation Server Using YaST” (page 246).

3.2.1.4 Starting the Installation

To start the installation, reboot the system. Right-click the profile name, select *Activate*, and press *OK* in the following dialog.

Use the screen console or connect to a serial console as described in the IBM documentation. One simple way to start a serial console is to open a VTerm while activating the partition. To do this, activate *Open a terminal window or console session* in the *Activate Logical Partition* dialog.

Enter the system firmware by pressing F1 or 1 when using a serial console or a virtual console during the system check when the system is rebooted:

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
```

```
1 = SMS Menu                      5 = Default Boot List
8 = Open Firmware Prompt          6 = Stored Boot List
```

```
memory      keyboard    network    scsi      speaker
```

Press F1 or 1 while the SCSI devices are checked. Select 5. *Select Boot Options* to enter the boot options dialog:

```
Version SF220_004
SMS 1.5 (c) Copyright IBM Corp. 2000,2003 All rights reserved.
-----
Main Menu
1.  Select Language
2.  Setup Remote IPL (Initial Program Load)
3.  Change SCSI Settings
4.  Select Console
5.  Select Boot Options
```

Navigation Keys:

X = eXit System Management Services

Type the number of the menu item and press Enter or select Navigation
Key:5

Select *1. Select Install/Boot Device* to set the *Install Device*. Go to *7. List all Devices* to see the list of available devices:

Version SF220_011

SMS 1.5 (c) Copyright IBM Corp. 2000,2003 All rights reserved.

Select Device

Device Number	Current Position	Device Name
1.	-	Virtual Ethernet (loc=U9111.520.10D3CCC-V1-C3-T1)
2.	-	Ethernet (loc=U787A.001.DNZ00XG-P1-T5)
3.	-	Ethernet (loc=U787A.001.DNZ00XG-P1-T6)
4.	-	IDE CD-ROM (loc=U787A.001.DNZ00XG-P4-D3)
5.	1	SCSI 73407 MB Harddisk (loc=U787A.001.DNZ00XG-P1-T10-L8-L0)

Navigation keys:

M = return to Main Menu

ESC key = return to previous screen X = eXit System Management Services

Type the number of the menu item and press Enter or select Navigation Key:

3.2.1.5 Booting from the CD-ROM Drive

Select the CD-ROM drive (4 in this example):

SMS 1.5 (c) Copyright IBM Corp. 2000,2003 All rights reserved.

Select Task

IDE CD-ROM

(loc=U787A.001.DNZ00XG-P4-D3)

1. Information
2. Normal Mode Boot
3. Service Mode Boot

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

Choose *2. Normal Mode Boot* to install from this device. On the next screen, confirm with *1. Yes* to exit *System Management Services* and boot from the device.

The system reads from the CD-ROM drive and the yaboot utility starts:

```
Welcome to SuSE:SLE-11:GA!

Type  "install"  to start the YaST installer on this CD/DVD
Type  "slp"      to start the YaST install via network
Type  "rescue"   to start the rescue system on this CD/DVD

Welcome to yaboot version 1.3.11.SuSE
Enter "help" to get some basic usage information
boot:
```

Type *install* and press Enter.

To read the installation data from a network install source rather than continuing the installation from the CD-ROM (see Section 3.2.1.3, “Assigning an Installation Device to an LPAR” (page 25)), append the option `manual` to the name of the kernel (`install`).

For an installation over VNC, append the parameters `vnc=1` and `vncpassword=password` to the name of the kernel (`install`). Read more about VNC in Section 14.1.1, “Simple Remote Installation via VNC—Static Network Configuration” (page 238).

3.2.1.6 Booting from the Network Source

Select an ethernet device that has access to the installation source (2 in this example).

3.2.1.7 Additional Steps

Proceed as described in Chapter 6, *Installation with YaST* (page 93) to begin installing the software with linuxrc and YaST.

3.2.2 Preparing for Installation on IBM pSeries Models

This section covers the preparatory steps for installing SUSE® Linux Enterprise Server on pSeries systems. It explains the installation from a built-in CD-ROM drive or a network source.

3.2.2.1 Special Features of IBM pSeries p630, p655, p670, and p690

IBM p630, p655, p670, and p690 systems offer the possibility of statically partitioning the system similarly to eServer p5/System p5 (which is described in Section 3.2.1, “Preparing for Installation on IBM eServer p5, System p, and OpenPower Models” (page 24)). This enables the concurrent operation of up to 16 operating systems on one machine. These operating systems are installed in *LPARs* (logical partitions). One or several of these partitions can contain a SUSE Linux Enterprise Server environment.

To prepare an LPAR for SUSE Linux Enterprise Server, first configure the system over the *HMC*. Refer to the Redbook *IBM eServer pSeries 690 System Handbook* (SG24-7040-00) for details (<http://www.redbooks.ibm.com/redbooks/SG247040/>).

Important notes regarding the configuration:

- The recommended maximum number of processors for a SUSE Linux Enterprise Server LPAR is eight, because the kernel can only manage eight processors effectively.
- For the installation, select *SMS* as the boot mode for the respective partition.
- The *HMC* terminal used for the input during the installation is a VT320 emulation. This emulation can lead to strange results with some applications. If possible, use an XTerm for communicating with the LPAR.

3.2.2.2 Hard Disk Space

Make sure that you have sufficient hard disk space for installing SUSE Linux Enterprise Server. The use of a separate hard disk is recommended.

SUSE Linux also supports installing to Fibre Channel–attached storage. Before beginning installation, the Fibre Channel Host Bus Adapter (FCHBA), SAN fabric, and storage system must each be configured to provide access from the FCHBA through the SAN Fabric to target logical units (LUNs) on the storage system.

SAN storage devices, if properly configured, are listed among existing hard disks on your system. *Create Custom Partitioning Setup* opens the dialog, as described in Section 15.1, “Using the YaST Partitioner” (page 275).


3.2.2.3 Setting Up the Installation Source

If you plan to install from CD-ROM, insert CD1 in the drive. In LPAR mode, the partition to install must have the CD-ROM in its partition profile. Create a network installation source if SUSE Linux Enterprise Server needs to be installed over a number of partitions. This eliminates the need to change CDs during installation. The same source can also be used for concurrent installation of various systems. The configuration of the network installation source is described in Section 14.2.1, “Setting Up an Installation Server Using YaST” (page 246).

3.2.2.4 Starting the Installation


To start the installation, reboot the system. Then enter the system firmware by pressing F1 or 1 when using the serial console during the system check when the system is rebooted. See Figure 3.3, “Entering the System Firmware” (page 33).

Figure 3.3: *Entering the System Firmware*



Press F1 or 1 while the SCSI devices are checked. Select 6 *MultiBoot* to enter the *Multiboot* dialog. See Figure 3.4, “Multiboot Dialog” (page 33).

Figure 3.4: *Multiboot Dialog*



Select 3 to set the install device. A list of available devices is displayed. See Figure 3.5, “Installing the Operating System” (page 34).

Figure 3.5: *Installing the Operating System*

```
InstallOperating System

Device Device
NumberName
1      Diskette
2      SCSI Tape id= 0 ( slot= 50322f5a )
3      SCSI CD-ROM id= 1 ( slot= 50322f5a )
4      Ethernet ( Integrated )
5      SysKonnnect PC IFDD IAdapter ( slot= 4 )
6      Ethernet ( slot= 2 )
7      None
```

3.2.2.5 Booting from the CD-ROM Drive

Select the respective CD-ROM drive (3 in this example). The system reads from the CD-ROM drive and displays the *identstring*.

```
->1    SuSE:SLE-11:GA<-
```

After you select 1, the yaboot utility is started.

```
Welcome to SuSE:SLE-11:GA!

Type "install"  to start the YaST installer on this CD/DVD
Type "slp"      to start the YaST install via network
Type "rescue"   to start the rescue system on this CD/DVD
```

Type *install* and press Enter. Alternatively, just press Enter to start the installer, the default option.

To install from a network source (see Section 3.2.2.3, “Setting Up the Installation Source” (page 32)), append `manual` to the kernel to `install`. For an installation over VNC, append the parameters `vnc=1` and `vncpassword=password` to `install`. Read more about VNC in Section 14.1.1, “Simple Remote Installation via VNC—Static Network Configuration” (page 238).

In LPAR mode, the partition to install must have the CD-ROM in its partition profile.

3.2.2.6 Booting from the Network Source

Select an ethernet device that has access to the installation source (6 in this example).

3.2.2.7 Additional Steps

Proceed as described in Chapter 6, *Installation with YaST* (page 93) to begin installing the software with `linuxrc` and YaST.

3.2.3 Preparing an Installation on IBM JSxx BladeCenter

This section describes the preparatory steps for the installation of SUSE® Linux Enterprise Server on JSxx Blades. It covers installation using the CD-ROM drive of the BladeCenter and using the network.

3.2.3.1 Creating a Network Installation Source

Create a network installation source if SUSE Linux Enterprise Server needs to be installed over a number of partitions. This provides the advantage of no CDs needing to be changed during installation. The same source can also be used for the concurrent installation of various systems. Configuration of a network installation source is described in Section 14.2.1, “Setting Up an Installation Server Using YaST” (page 246).

3.2.3.2 Hard Disk Storage Space

Ensure that enough hard disk storage space is available for the installation of SUSE Linux Enterprise Server. It is recommended you use a dedicated hard disk.

3.2.3.3 Preparing the System for Boot

Preparing to Boot from the CD-ROM Drive

Perform the steps described in this section if an installation from CD-ROM is desired.

Assign the CD-ROM drive to the Blade chosen for installation by connecting (with a Web browser) to a BladeCenter Management Module, then logging in. After login, select the function *Remote Control* in the menu *Blade Tasks* then activate *Start Remote Control*. Assign the CD-ROM drive to the desired blade in the menu *Change Media Tray Owner* of the new window.

Set up the CD-ROM drive as a boot device. Do this by selecting *Blade Tasks* then *Configuration* while in the BladeCenter Management Module. Select the JSxx Blade in the section *Boot Sequence*. Set the entry for *1st Device* on the page for *Blade Boot Sequence* to *CDROM*.

Put CD 1 in the CD-ROM drive and restart the blade.

Preparing to Boot from the Network

Perform the steps as described in this section if an installation over the network is desired.

Connect to the BladeCenter Management Module using a Web browser and log in. Set the boot device to the network by accessing the *Configuration* menu from the *Blade Tasks* page. Select the JSxx Blade in the section *Boot Sequence* and set *1st Boot Device* to *Network — BOOTP* on *Blade Boot Sequence*.

Rebooting and Connecting to the Console of the JSxx Blade

Reboot the JSxx Blade from the item *Power/Restart* of the *Blade Tasks* menu in the BladeCenter Management Module. A table appears, showing the power status of the

blades in the *Pwr* column. Mark the check box of the desired blade and restart it with *Power On Blade*.

Connect to the BladeCenter with the command `telnet bladecenter` and log in.

```
username: user
password: *****
system>
```

The command `env -T system:blade[bay number]` determines for which JSxx Blade the subsequent commands are intended. The blades installed in the BladeCenter are listed by calling `list -l 3`.

```
system> list -l 3
system
    mm[1]      primary
    power[1]
    power[2]
    power[3]
    power[4]
    blower[1]
    blower[2]
    switch[1]
    switch[3]
    blade[1]
        sp
        cpu[1]
        cpu[2]
    blade[3]
        sp
    blade[4]
        sp
    blade[6]
        sp
    blade[8]
        sp
        cpu[1]
        cpu[2]
    blade[9]
        sp
        cpu[1]
        cpu[2]
    blade[10]
        sp
    blade[11]
        sp
    blade[13]
        sp
    mt
```

```
system>
```

The *command target* is then determined. To work, for example, with blade number 9, enter `env -T system:blade[9]`. Connect with the console of the JSxx Blade over *Serial over LAN (SOL)* with the command `console`.

```
system> env -T system:blade[9]
OK
system:blade[9]> console
```

Starting the Installation

The SUSE Linux Enterprise Server boot loader starts after the system check has completed.

```
Welcome to SuSE:SLE-11:GA!

Type  "install"  to start the YaST installer on this CD/DVD
Type  "slp"      to start the YaST install via network
Type  "rescue"   to start the rescue system on this CD/DVD

Welcome to yaboot version 1.3.11.SuSE
Enter "help" to get some basic usage information
boot:
```

Select *install* from the menu and press Enter.

In the case of an installation over VNC, append the parameters `vnc=1` and `vncpassword=password` to the command line for the kernel (install).

Additional Steps

Proceed as described in Chapter 6, *Installation with YaST* (page 93) to begin installing the software with `linuxrc` and YaST.

Installation on IBM System z

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

4.1 General Information and Requirements

This section gives basic information about the system requirements (like supported hardware), level of MicroCode, and software. It also covers the different installation types, how to do an IPL for the first installation, and information about the IOCDS.

4.1.1 System Requirements

This section provides a list of hardware for IBM System z supported by SUSE Linux Enterprise Server. Next, the level of the MicroCode (MCL) used in your IBM System z system, which is very important for the installation, is covered. Additional software to install and use for installation is mentioned at the end of this section.

4.1.1.1 Hardware

SUSE Linux Enterprise Server has run successfully on the following platforms:

- IBM Series z9 (z9-EC) 2094
- IBM Series z9 (z9-BC) 2096
- IBM Series z10 (z10-EC) 2097
- IBM Series z10 (z10-BC) 2098
- IBM zEnterprise System z196 2817
- IBM zEnterprise System z114 2818
- IBM zEnterprise EC12 (zEC12) 2827

Memory Requirements

Different installation methods have different memory requirements during installation. After installation is completed, the system administrator may reduce memory to the desired size. SUSE recommends using:

768 MB	For installation under z/VM.
1 GB	For installation under LPAR.

NOTE: Memory Requirements with Remote Installation Sources

For installation from NFS, FTP, or SMB installation sources or whenever VNC is used, 512MB of memory is required as a minimum. Otherwise, the installation attempt is likely to fail. Further note that the number of devices visible to the z/VM guest or LPAR image affects memory requirements. Installation with literally hundreds of accessible devices (even if unused for the installation) may require more memory.

Disk Space Requirements

The disk requirements depend largely on the installation. Commonly, you need more space than the installation software itself needs to have a system that works properly. Minimal requirements for different selections are:

2.6 GB	Default Installation
--------	----------------------

3.6 GB+

Recommended (this is with graphical desktop, development packages and Java).

Network Connection

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

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

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

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

IPL Options

For an LPAR installation, the *Load from CD-ROM or Server* option is the preferred way to IPL the installation kernel and initrd (initial RAM disk). If this option is not available and you cannot use z/VM for installing the system, you need to IPL from a channel attached tape with the tapeipl kernel, the parmfile, and the initrd. Thus, you need access to a tape unit (3480, 3490, or 3590, for example).

4.1.1.2 MicroCode Level, APARs, and Fixes

Documentation about restrictions and requirements for this release of SUSE Linux Enterprise Server can be found on IBM developerWorks at http://www.ibm.com/developerworks/linux/linux390/documentation_suse.html. It is recommended always to use the highest service level available. Contact your IBM support for minimum requirements.

z/VM

z/VM 5.4

z/VM 6.2

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

4.1.1.3 Software

To install 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 installing via SMB on these machines is generally recommended.

To connect to the SUSE Linux Enterprise Server installation system, one of the following methods is required:

SSH with Terminal Emulation (xterm compatible)

SSH is a standard Unix tool that should be present on any Unix or Linux system. For Windows, there is an SSH client called Putty. It is free to use and is available from <http://www.chiark.greenend.org.uk/~sgtatham/putty/>.

VNC Client

For Linux, a VNC client called vncviewer is included in SUSE Linux Enterprise Server as part of the `tightvnc` package. For Windows, `tightvnc` is also available. Download it from <http://www.tightvnc.com/>. Alternatively, use the VNC Java client and a Java-enabled Web browser.

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 Macintosh. Some of them can be downloaded as free trial versions. A trial version of the Mocha X Server from MochaSoft can be obtained at <http://www.mochasoft.dk/freeware/x11.htm>.

TIP: Additional Information

Consult the `README` located in the root directory of DVD 1 of your SUSE Linux Enterprise Server before installing SUSE Linux Enterprise Server

on IBM System z. This file completes the documentation presented in this book.

4.1.2 Installation Types

This section gives an overview of the different types of installation possible with SUSE Linux Enterprise Server for IBM System z. Basically, these two types are given:

LPAR

Installation of SUSE Linux Enterprise Server using a logical partition (LPAR).

VM (z/VM)

Installation of SUSE Linux Enterprise Server as a guest operating system within z/VM.

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

4.1.2.1 LPAR

If you install SUSE Linux Enterprise Server for IBM System z into a separate logical partition (LPAR), allow SUSE Linux Enterprise Server to use a special part of the physical memory in your system. Also decide how many processors are used by SUSE Linux Enterprise Server. In this mode, you can run different operating systems simultaneously on your IBM System z system.

4.1.2.2 z/VM

Running SUSE Linux Enterprise Server for IBM System 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 are able to run hundreds of Linux instances simultaneously.

4.1.3 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 channel-attached tape, VM reader, and load from CD-ROM or server options are discussed. Installing the software packages, which is done over the network, does not require the IPL medium.

4.1.3.1 ESCON or FICON Attached Tape

IPLing from a channel-attached tape is possible on all systems connected to a tape library. The only prerequisite is that the LPAR in which to install (or allowing z/VM to run) is allowed to access the tape unit. For this, the `IODEVICE` statement in the IOCDS must have the attribute `SHARED` or `PART=<LPARName>`.

4.1.3.2 VM Reader

To IPL from a VM reader, transfer the necessary files into the reader first. Then multiple IPLs are easily done. This is the preferred way on z/VM. 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.

4.1.3.3 Load from CD/DVD-ROM or Server

For IPLing into an LPAR, it is possible to either 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 where the data is to be copied. For SUSE Linux Enterprise Server this file is called `suse.ins` and located in the root directory of the file system on the DVD 1.

In the left navigation pane of the HMC expand *Systems Management and Servers* 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 displayed in the upper content area on the right. In the Tasks area, expand *Recovery* and click *Load from CD-ROM, DVD, 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. In case the `suse.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 `suse.ins` entry. Start the installation with *OK*.

4.1.3.4 Load from SCSI-Attached DVD

To IPL from a SCSI DVD, you need access to an FCP adapter connected to a DVD drive. You need values like the WWPN and LUN from the SCSI drive. For details, see Section “IPL from FCP-Attached SCSI DVD” (page 57).

4.1.3.5 Load from the Network with zPXE

IPLing from the Network with zPXE requires a Cobbler server providing the kernel, RAM disk and a `parmfile`. zPXE is only available on z/VM and is initiated by running the ZPXE EXEC script. See Section 4.2.1.3, “Using a Cobbler Server for zPXE” (page 50) for details.

4.1.4 The IOCDs

This section provides some necessary information about the IOCDs and how to customize some settings for sharing network cards or DASDs among several LPARs. In the IOCDs, the `chpid` and types of the devices connected to the IBM System z are defined. The resources can be dedicated or shared among LPARs.

WARNING: Sharing Devices (DASD)

Do not share writable DASD among LPARs because this might result in data loss. Consider the definition of the necessary resources in advance when planning the setup for SUSE Linux Enterprise Server on IBM System z.

This example shows how to dedicate a DASD to one specific LPAR. This LPAR is referred to as `LPAR1`.

Example 4.1: *Dedicating DASD to One LPAR*

```
CHPID PATH=FD,TYPE=DSD,SHARED
CNTLUNIT CUNUMBR=FD00,PATH=FD,UNITADD=((00,256)),UNIT=3990-2
IODEVICE ADDRESS=(FD03,1),CUNUMBR=FD00,UNIT=3990,PART=LPAR1
```

To share a DASD among LPARs, delete the `PART=LPAR1` part in the IOCDS definition. This might be useful for high availability reasons or for sharing data among LPARs read-only.

Several Linux systems can use the same network device if you share it among LPARs or z/VM guests. This reduces the number of network devices that must be provided to the Linux system. On the other hand, you might provide more than one network device to one Linux system to make it more available in case one connection fails.

Network cards like OSA-Express can be used in two different modes. These modes are known as QDIO and non-QDIO mode. Define these modes in the IOCDS by using the `TYPE` statement. QDIO mode is much faster than non-QDIO mode, but uses three device addresses instead of two in non-QDIO. Consider the limited number of device addresses when planning the setup of your IBM System z Linux environment.

Example 4.2: *Sharing OSA Express Card among LPARs (non-qdio) on z9*

```
CHPID PATH=(FE),SHARED,PARTITION=((LPAR1,LPAR2)),TYPE=OSE
CNTLUNIT CUNUMBR=FE00,PATH=(FE),UNIT=OSA
IODEVICE ADDRESS=(FE00,016),CUNUMBR=(FE00),UNIT=OSA
IODEVICE ADDRESS=(FEFE,001),CUNUMBR=(FE00),UNIT=OSAD
```

Example 4.3: *Sharing OSA Express Card among LPARs (qdio) on z9*

```
CHPID PATH=(FE),SHARED,PARTITION=((LPAR1,LPAR2)),TYPE=OSD
CNTLUNIT CUNUMBR=FE00,PATH=(FE),UNIT=OSA
IODEVICE ADDRESS=(FE00,016),CUNUMBR=(FE00),UNIT=OSA
IODEVICE ADDRESS=(FEFE,001),CUNUMBR=(FE00),UNIT=OSAD
```

4.2 Preparing for Installation

In this section, learn 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. Also find out about network configuration and network installation.

4.2.1 Making the Installation Data Available

This section provides detailed information about making the SUSE Linux Enterprise Server IBM System 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 also use the Windows network (including the SMB protocol) to install SUSE Linux Enterprise Server on your IBM System z system.

TIP: IPL from DVD

Since Service Pack 1 of SUSE Linux Enterprise Server Version 10, 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 install from hard disk by putting the content of the DVD to a partition on a DASD.

4.2.1.1 Using a Linux Workstation or SUSE Linux Enterprise Server DVD

If you have a Linux workstation running in your computer environment, use the workstation to provide the installation data to the IBM System z installation process by NFS or FTP. If the Linux workstation runs under SUSE Linux Enterprise Server, you can set up an installation server (NFS or FTP) using the YaST *Installation Server* module as described in Section 14.2.1, “Setting Up an Installation Server Using YaST” (page 246).

Over NFS

Use NFS (network file system) to make the installation media available.

IMPORTANT: Exporting Mounted Devices with NFS

Exporting the file system root (/) does not imply the export of mounted devices, such as DVD. Explicitly name the mount point in `/etc/exports`:

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

After changing this file, restart the NFS server with the command
`rcnfsserver restart`.

Over FTP

Setting up an FTP server on a Linux system involves the installation of the server software itself, such as `wuftp` or `proftpd`, as well as other possible configuration tasks. Using YaST, the installation step is straightforward: select the package to install and start the installation. Skip the configuration of the FTP server if no anonymous FTP should be used for the installation. Instead, use an FTP login with a valid username and password. You might want to create a user account for this task only. The FTP daemon does not need to be started by hand. It can be started by `inetd` if an FTP connection is requested. To activate the new settings, enter `rcinetd restart` or `rcxinetd restart`.

SUSE Linux Enterprise Server on DVD

DVD1 of the SUSE Linux Enterprise Server for IBM System z contains a bootable Linux image for Intel-based workstations as well as an image for System z.

For Intel-based workstations, boot from this DVD, answer the questions regarding your 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 takes some Linux and networking experience, because you need to set up the networking of the workstation manually.

For System z, IPL your LPAR/VM guest from this DVD as described in Section “IPL from FCP-Attached SCSI DVD” (page 57). After entering your network parameters, the installation system treats the DVD as the source of installation data. Because System z cannot have an X11-capable terminal attached directly, choose between VNC or SSH installation. SSH also provides a graphical installation by tunneling the X connection through SSH with `ssh -X`.

4.2.1.2 Using a Microsoft Windows Workstation

If there is a Microsoft Windows workstation available in your network, use this computer to make the installation media available. The easiest way to do this is to use the SMB protocol, already included in the Windows operating system. Be 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. Another option is to use FTP. This also requires some third-party software for Windows.

With SMB

To make the installation media available with SMB, just insert the SUSE Linux Enterprise Server DVD 1 into the DVD drive of the Windows workstation. Then create a new share using the DVD-ROM 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 username 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).

With NFS

Refer to the documentation provided with the third party product that enables NFS server services for your Windows workstation. The DVD-ROM drive containing the SUSE Linux Enterprise Server DVDs must be in the available NFS path.

With FTP

Refer to the documentation provided with the third party product that is enabling FTP server services on your Windows workstation. The DVD-ROM drive containing the SUSE Linux Enterprise Server DVDs must be in the available FTP path.

The FTP server that is bundled with some Microsoft Windows releases implements only a subset of the FTP command set and is not suitable for providing the installation data. However, other products (such as the FTP server that is part of Hummingbird Exceed or WAR-FTPD) have been reported as functional.

Using an FCP-Attached SCSI DVD Drive

After you IPLed from the SCSI DVD as described in Section 4.1.3.4, “Load from SCSI-Attached DVD” (page 45), the installation system uses the DVD as the installation medium. In that 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 by VNC or by X tunneled through SSH.

4.2.1.3 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 four steps:

- Importing the Installation Data
- Adding a Distribution
- Adding Profiles
- Adding Systems

Importing the Installation Data

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

```
cobbler import --path=PATH❶ --name=IDENTIFIER❷ --arch=s390x
```

- ❶ Mount point of the installation data.
- ❷ A string identifying the imported product, for example “sles11_sp3_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.

Adding a Distribution

By adding a distribution, you tell 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 System z:

```
cobbler distro add --arch=s390x --breed=suse --name="IDENTIFIER"❶ \  
  --os-version=sles10❷ \  
  --initrd=/srv/www/cobbler/ks_mirror/IDENTIFIER/boot/s390x/initrd❸ \  
  --kernel=/srv/www/cobbler/ks_mirror/IDENTIFIER/boot/s390x/vmldr.ikr❹ \  
  --kopts="install=http://cobbler.example.com/cobbler/  
ks_mirror/IDENTIFIER"❺
```

- ❶ Custom identifier for the distribution, for example “SLES 11 SP3 System z”. Must be unique.
- ❷ Operating system identifier. Use `sles10`.
- ❸ Path to the initrd. The first part of the path (`/srv/www/cobbler/ks_mirror/IDENTIFIER/`) depends on the location where Cobbler imported the 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 where Cobbler imported the data and the subdirectory name you chose when importing the installation data.
- ❺ URL to the installation directory on the Cobbler server.

4.2.1.4 Adding Profiles

With a profile you can add additional options to a distribution—for example adding an AutoYaST file for an automated installation. You can specify multiple profiles per distribution; at least one must be created.

```
cobbler profile add  
  --name=PROFILENAME❶ --distro=DISTRIBUTION❷ --  
kickstart=PATH_TO_AUTOYAST_FILE❸
```

- ❶ Unique name for the profile.
- ❷ Distribution to which the profile should apply. You must use the string specified with `--name=IDENTIFIER` in the importing step here.

- ③ Specify the path to an AutoYaST file for an automated installation here. This parameter is optional.

4.2.1.5 Adding Systems

The last step that is required is to add systems to the Cobbler server. A system addition needs to be done for every System z guest that should boot via zPXE. Guests are identified via their z/VM user ID (in the following example, an ID called “LINUX01” is assumed). To add a system, run the following command:

```
cobbler system add --name=LINUX01 --hostname=linux01.example.com \  
--ip=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"
```

With the `--kopts` option you can specify the kernel and installation parameters you would normally specify in the `parmfile`. The parameters are entered as a space-separated list in the form of `PARAMETER1=VALUE1 PARAMETER2=VALUE2`. The installer will prompt you for missing parameters. For a completely automated installation you need to specify all parameters for networking, DASDs and provide an AutoYaST file. The following shows 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 \  
OsaMedium=eth \  
Layer2=0 \  
PortNo=0 \  
ReadChannel=0.0.0600 \  
WriteChannel=0.0601 \  
DataChannel=0.0.0602 \  
Portname=DT70 \  
DASD=600"
```

4.2.2 Installation Types

This section provides information about which steps must be performed to install SUSE Linux Enterprise Server for each of the installation modes and where to find the appropriate information. After the preparations mentioned in the previous chapters have been accomplished, follow the installation overview of the desired installation mode to install SUSE Linux Enterprise Server on your system.

As described in Section 4.2.1, “Making the Installation Data Available” (page 46), there are two different installation modes for Linux on IBM System z:

- LPAR Installation
- z/VM Installation

Procedure 4.1: Overview of LPAR Installation

- 1** Prepare the devices needed for installation. See Section 4.2.3.1, “LPAR Installation” (page 54).
- 2** IPL the installation system. See Section 4.2.4.1, “LPAR Installation” (page 56).
- 3** Configure the network. See Section 4.2.5, “Network Configuration” (page 61).
- 4** Connect to the SUSE Linux Enterprise Server installation system. See Section 4.2.6, “Connecting to the SUSE Linux Enterprise Server Installation System” (page 64).
- 5** Start the installation using YaST and IPL the installed system. See Chapter 6, *Installation with YaST* (page 93).

Procedure 4.2: Installation Overview of z/VM Installation

- 1** Prepare the devices needed for installation. See Section 4.2.3.2, “z/VM Installation” (page 55).
- 2** IPL the installation system. See Section 4.2.4.2, “z/VM Installation” (page 58).
- 3** Configure the network. See Section 4.2.5.1, “z/VM Installation” (page 62).

- 4 Connect to the SUSE Linux Enterprise Server installation system. See Section 4.2.6, “Connecting to the SUSE Linux Enterprise Server Installation System” (page 64).
- 5 Start the installation using YaST and IPL the installed system. See Chapter 6, *Installation with YaST* (page 93).

4.2.3 Preparing the IPL of the SUSE Linux Enterprise Server Installation System

4.2.3.1 LPAR Installation

Configure your IBM System z system to start in ESA/S390 or LINUX-only mode with an appropriate activation profile and IOCDS. Consult IBM documentation for more on how to achieve this.

IOCDS: Attaching and Configuring Devices

A SUSE Linux Enterprise Server installation needs at least two devices: a DASD and a network connection device. For an IPL from tape, a tape device should also be accessible. Devices are configured and attached to an LPAR in the IOCDS (input output configuration data set). This example defines one DASD, one OSA-2 network device, and a tape device for LPAR Z1. For further information about how to set up the IOCDS for Linux, refer to your machine's IBM hardware documentation.

Example 4.4: *An Example IOCDS*

```
CHPID PATH=(CSS(0),FD),PCHID=120,TYPE=FC
CHPID PATH=(CSS(0),FE),PCHID=320,TYPE=OSD
CHPID PATH=(CSS(0),10),PCHID=3A0,TYPE=CNC

CNTLUNIT CUNUMBR=FD00,PATH=((CSS(0),FD)),UNITADD=((00,1)),UNIT=2105
IODEVICE ADDRESS=(FD00,1),CUNUMBR=(FD00),UNIT=3390B,UNITADD=00

CNTLUNIT CUNUMBR=FE20,PATH=((CSS(0),FE)),UNIT=OSA
IODEVICE ADDRESS=(FE20,1),CUNUMBR=(FE20),UNIT=OSA
IODEVICE ADDRESS=(FEFE,1),CUNUMBR=(FE20),UNIT=OSAD

CNTLUNIT CUNUMBR=100A,PATH=((CSS(0),10)),UNIT=3480,UNITADD=((0A,1))
IODEVICE ADDRESS=(100A,1),CUNUMBR=(100A),UNIT=3480,UNITADD=00
```

Proceed with Section 4.2.4.1, “LPAR Installation” (page 56).

4.2.3.2 z/VM Installation

Adding a Linux Guest

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`, 256 MB of memory (extendable up to 1024 MB), 32 MB of expanded RAM (XSTORE), some 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 suits the needs of your preferred installation type. See Section “Memory Requirements” (page 40). To set the memory size to 512 MB, use the command `CP DEFINE STORAGE 512M`. After the installation has finished, reset the memory size to the desired value.

Example 4.5: *Configuration of a z/VM Directory*

```
USER LINUX1 LINPWD 256M 1024M G
*
* LINUX1
*
* This VM Linux guest has two CPUs defined.

CPU 01 CPUID 111111
CPU 02 CPUID 111222
IPL CMS PARM AUTO CR
IUCV ANY
IUCV ALLOW
MACH ESA 10
OPTION MAINTCCW RMCHINFO
SHARE RELATIVE 2000
XSTORE 32M
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 TWO4ME THR4ME
MDISK 150 3390 0052 0200 DASD40 MR ONE4ME TWO4ME THR4ME
MDISK 151 3390 0253 2800 DASD40 MR ONE4ME TWO4ME 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.

Now add (as the user MAINT) 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)
```

Assign the directories to the guest with `DISKMAP USER DIRECT A`. 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 to set up this user.

Proceed with Section 4.2.4.2, “z/VM Installation” (page 58).

4.2.4 IPLing the SUSE Linux Enterprise Server Installation System

4.2.4.1 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.

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. In the list of options that appears, choose the default selection. *Operating system messages* should now show the kernel boot messages.

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.

IPL from ESCON or FICON Attached Tape

If you cannot IPL from DVD, create a channel attached tape from which to IPL the SUSE Linux Enterprise Server installation image. Use the *LOAD* button in the SE or HMC with the tape device address as the load address to IPL the SUSE Linux Enterprise Server installation system.

There are many ways to create an IPLable tape. One is to copy the files:

```
/boot/s390x/tapeipl.ikr  
/boot/s390x/parmfile  
/boot/s390x/initrd
```

as binary files from DVD 1 (for example, using FTP from a Linux workstation).

Name them

```
SLES11 IMAGE  
SLES11 PARM  
SLES11 INITRD
```

and write them onto a tape with the REXX from the example.

IMPORTANT: Transferring Binaries using FTP

Do not upload the files as `fixed 80`. Store them as `fixed 1024`. Use the FTP command `locsite fix 1024`.

Example 4.6: *REXX Script to Create an IPLable Tape*

```
'REWIND 181'  
'FILEDEF IN1 DISK' SLES11 IMAGE A  
'FILEDEF IN2 DISK' SLES11 PARM A  
'FILEDEF IN3 DISK' SLES11 INITRD A  
'FILEDEF OUT TAP1 (RECFM F BLOCK 1024 LRECL 1024 PERM'  
say 'Writing: ' left(file1,23)  
'MOVEFILE IN1 OUT'  
say 'Writing: ' left(file2,23)  
'MOVEFILE IN2 OUT'
```

```

say 'Writing: ' left(file3,23)
'MOVEFILE IN3 OUT'
say 'Done.'
'REWIND 181'
exit

```

The tape in this script is attached as 181. Adjust the script to your needs.

4.2.4.2 z/VM Installation

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

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` on DVD 1 of the SUSE Linux Enterprise Server for IBM System z available by FTP within your network. From this directory, get the files `vmrdr.ikr`, `initrd`, `parmfile`, and `sles11.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 `vmrdr.ikr` (the Linux kernel) and `initrd` (the installation image) as binary files, so use the `binary` transfer mode. `parmfile` and `sles11.exec` need to be transferred in ASCII mode.

The example shows the steps necessary. In this example, the required files are accessible from an FTP server at the IP address `192.168.0.3` and the login is `lininst`. It may differ for your network.

Example 4.7: *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/vmrdr.ikr sles11.image
200 PORT Command successful
150 Opening BINARY mode data connection for /media/dvd1/boot/s390x/vmrdr.ikr
(6757376 bytes)
226 Transfer complete.
6757376 bytes transferred in 8.826 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
get /media/dvd1/boot/s390x/initrd sles11.initrd
200 PORT Command successful
150 Opening BINARY mode data connection for /media/dvd1/boot/s390x/initrd
(12654815 bytes)
226 Transfer complete.
12194534 bytes transferred in 16.520 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
ascii
200 Type set to A
Command:
get /media/dvd1/boot/s390x/parmfile sles11.parmfile
150 Opening ASCII mode data connection for /media/dvd1/boot/s390x/parmfile
(71 bytes)
226 Transfer complete.
71 bytes transferred in 0.092 seconds.
Transfer rate 0.71 Kbytes/sec.
Command:
get /media/dvd1/boot/s390x/sles11.exec sles11.exec
150 Opening ASCII mode data connection for /media/dvd1/boot/s390x/
sles11.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 `sles11.exec` you just downloaded to IPL the Linux installation system. This script loads the kernel, `parmfile`, and the initial RAM disk into the reader for IPL.

Example 4.8: SLES11 EXEC

```

/* REXX LOAD EXEC FOR SUSE LINUX S/390 VM GUESTS          */
/* LOADS SUSE LINUX S/390 FILES INTO READER              */
SAY ''

```

```

SAY 'LOADING SLES11 FILES INTO READER...'
'CP CLOSE RDR'
'PURGE RDR ALL'
'SPOOL PUNCH * RDR'
'PUNCH SLES11 IMAGE A (NOH'
'PUNCH SLES11 PARMFILE A (NOH'
'PUNCH SLES11 INITRD A (NOH'
'I 00C'

```

With this script you can IPL the SUSE Linux Enterprise Server installation system with the command `sles11`. The Linux kernel then starts and prints its boot messages.

To continue the installation, proceed to Section 4.2.5.1, “z/VM Installation” (page 62).

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 4.2.5.1, “z/VM Installation” (page 62).

IPL from ESCON or FICON Attached tape

If you cannot IPL from a z/VM reader, create a channel attached tape from which to IPL the SUSE Linux Enterprise Server installation image. For instructions, refer to Section “IPL from ESCON or FICON Attached Tape” (page 57).

To continue the installation, proceed with Section 4.2.5.1, “z/VM Installation” (page 62).

IPL from a Cobbler Server with zPXE

To IPL from a Cobbler server with zPXE you need to transfer the `zpxe.exec` script via FTP from the Cobbler server to your z/VM guest. 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 4.7, “Transferring the Binaries via FTP” (page 58) for an example). The `zpxe.exec` script is available on the Cobbler server at `ftp://IP_OF_COBBLER_SERVER/zSERIES_INSTALLATION_DIRECTORY/boot/s390x/zpxe.exec`. The exact location of `zSERIES_INSTALLATION_DIRECTORY` depends on where you imported the installation data on the Cobbler server (see Section “Importing the Installation Data” (page 50) for details).

`zpxe.exec` is supposed to replace the `PROFILE EXEC` of your guest. Make a backup copy of the existing `PROFILE EXEC` and rename `ZPXE EXEC` to `PROFILE EXEC`. Alternatively call `ZPXE EXEC` from the existing `PROFILE EXEC` by using a new line with the following content: '`ZPXE EXEC`'.

The last step is to create a config file, `ZPXE CONF`, telling `ZPXE EXEC` 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
IPLDISK 600
```

On the next log in to your z/VM guest, the Cobbler server will be connected. 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:

```
cobbler system edit --name ID❶ --netboot-enabled 1❷ --profile PROFILENAME❸
```

- ❶ z/VM user ID.
- ❷ Enable IPLing from the network.
- ❸ Name of an existing profile, see Section 4.2.1.4, “Adding Profiles” (page 51).

4.2.5 Network Configuration

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

First, choose *Start Installation or System* in the `linuxrc` main menu then *Start Installation or Update* to start the installation process. Select *Network* as your installation medium then select the type of network protocol you will

be utilizing for the installation. Section 4.2.1, “Making the Installation Data Available” (page 46) 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.

Now set up the network device over which to receive the installation data: *OSA-2* or *OSA Express* or *HiperSockets*. The following network adapters are still available and usable, but no longer supported: CTC, ESCON, IUCV. Next, choose the CCW bus interface and the physical medium (*Ethernet*). As a result, the respective driver is installed and you see the corresponding kernel messages.

Proceeding with the installation, linuxrc displays a list of potential usable read, write, and, if applicable, data channels. After entering the addresses for each channel, you may also need to enter additional information, such as the port name for OSA ethernet cards.

Next, decide 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, you probably want to say *NO* here. When you do so, you are prompted for the networking parameters of your installation network device:

- The IP address of the system to install
- The corresponding netmask
- The IP address of a gateway to reach the server
- The IP address of your domain name server (DNS)

When using an OSA Express Network Card you are now prompted for a *relative port number*. This was added to support the new 2 port OSA Express 3 Network devices. If you are not using an OSA Express 3 device, please enter 0. OSA Express cards also have the option of running in an “OSI layer 2 support” mode or using 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, please 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.

4.2.5.1 z/VM Installation

After the kernel has completed its start-up routines, answer a few questions regarding the network setup. First, select the type of network connection to use: OSA Express or HiperSockets. In this example installation, OSA Express is used.

The system now displays a possible OSA configuration. Choose first whether to use QDIO or LCS OSA. Next, choose the physical medium to use and enter the device addresses. If you prefer another setup, enter the device address of the OSA read channel (0.0.0700 in this example) then the one of the OSA write channel (0.0.0701) and the OSA control channel (0.0.0702). After entering the channels, insert the name of the port to which the OSA card is connected.

SUSE Linux Enterprise Server now tries to load the network module by building a parameter line with the information provided and displaying all loaded modules. Loading was successful if you get an output like:

Example 4.9: *Network Device Driver Parameters*

```
(portname YSW2)
(Port 0)
qdio: 0.0.0702 OSA on SC 3 using AI:1 QEBSM:0 PCI:1 TDD:1 SIGA:RW AO
geth.736dae: 0.0.0700: Device is a Guest LAN QDIO card (level: V540)
with link type GuestLAN QDIO (portname: YSW2)
geth.47953b: 0.0.0700: Hardware IP fragmentation not supported on eth0
geth.066069: 0.0.0700: Inbound source MAC-address not supported on eth0
geth.d7fdb4: 0.0.0700: VLAN enabled
geth.e90c78: 0.0.0700: Multicast enabled
geth.5a9d02: 0.0.0700: IPV6 enabled
geth.184d8a: 0.0.0700: Broadcast enabled
geth.dac2aa: 0.0.0700: Using SW checksumming on eth0.
geth.9c4c89: 0.0.0700: Outbound TSO not supported on eth0
```

Next, enter your IP address, netmask, and default gateway. To install over iucv or ctc, enter additional information, like the peer address (for a point-to-point adapter) or the port name.

Finally, the IP address of the DNS server and the MTU size are requested. The MTU size should always match the one used by the network to which you are connecting.

Now a summary is displayed. Confirm if your input is correct. Before the network is started, enter a password that is valid only during the installation. After having IPLed the installed system, enter the *real* root password.

With all basic parameters set up, the network is started. Check the output of ifconfig, which should contain two entries: a loopback (lo) connection and one connection (eth0, ctc0, escon0, iucv0, or hsi0) with correct settings.

Example 4.10: *Example ifconfig*

```
/sbin/ifconfig eth0 :  
Link encap:Ethernet HWaddr 02:00:01:00:00:27  
inet addr:192.168.0.1 Bcast:192.168.0.255 Mask:255.255.255.0  
inet6 addr: fe80::200:100:100:27/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1492 Metric:1  
RX packets:0 errors:0 dropped:0 overruns:0 frame:0  
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:1000  
RX bytes:0 (0.0 Mb) TX bytes:0 (0.0 Mb)
```

4.2.6 Connecting to the SUSE Linux Enterprise Server Installation System

After setting up your network connection, `linuxrc` prompts for the details of the installation source chosen earlier in the process, for example, the server IP address and the directory in which the data is located.

Finally, `linuxrc` wants to know what type of display you want to use to control the installation procedure. Possible choices are `X11` (X Window System), `VNC` (Virtual Network Computing protocol), `SSH` (text mode or `X11` installation via Secure Shell), or `ASCII Console`.

When choosing the latter (`ASCII Console`), `YaST` will be started in text mode and you can perform the installation directly within your terminal. See Chapter 3, *YaST in Text Mode* (*Administration Guide*) for usage instructions.

NOTE: Terminal Emulation for ASCII Console

In order to be able to work with `YaST` in text mode, it needs to run in a terminal with `VT220/Linux` emulation (also referred to as `ASCII console`). You will not be able to use `YaST` in a `3270` terminal, for example.

4.2.6.1 Initiating the Installation for VNC

- 1 After the installation option `VNC` has been chosen, the `VNC` server starts. A short note displayed in the console provides information about which IP address and display number is needed for a connection with `vncviewer`. Alternatively,

a URL is given here for entry into your Java-enabled browser to connect to the installation system.

- 2 Start a VNC client application on your client system. Either use `vncviewer` or the VNC Java client and a Java-enabled Web browser.
- 3 Enter the IP address and the display number of the SUSE Linux Enterprise Server installation system when prompted to do so.

If you connect via a Java-enabled browser, enter a URL containing the IP address of the installation system and the appropriate port number in the format:

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

- 4 After the connection has been established, start installing SUSE Linux Enterprise Server with YaST.

4.2.6.2 Initiating the Installation for the X Window System

IMPORTANT: X Authentication Mechanism

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

- 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`. Then 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.

4.2.6.3 Initiating the Installation for SSH

To connect to an installation system with the name `earth` using SSH, execute `ssh -X earth`. If your workstation runs on Microsoft Windows, use the `ssh` and `telnet` client and terminal emulator `putty`, which is available from <http://www.chiark.greenend.org.uk/~sgtatham/putty/>. Set *Enable X11 forwarding* in `putty` under *Connection > SSH > X11*.

A login prompt appears. Enter `root` and log in with your password. Enter `yast2` to start YaST.

Proceed with the detailed description of the installation procedure that can be found in Chapter 6, *Installation with YaST* (page 93).

4.3 Network Connection Types

SUSE Linux Enterprise Server for IBM System z includes network drivers for OSA devices (ethernet, and gigabit ethernet) and HiperSockets. This chapter describes the configuration within the SUSE Linux Enterprise Server installation system.

WARNING: CTC, ESCON, and IUCV Interfaces No Longer Supported

CTC, ESCON, and IUCV interfaces are no longer officially supported. For compatibility reasons, they are still usable, but with the next release of SUSE Linux Enterprise Server the support of these interfaces will be dropped completely.

4.3.1 HiperSockets

Select your device from the list of network devices. Then enter the network device read channel number (such as 0.0.700), the write channel number (like 0.0.701) and the data channel number (like 0.0.702).

Example 4.11: *Supported Network Connection Types and Driver Parameters*

Choose the network device.

- 1) IBM parallel CTC Adapter (0.0.0600)
- 2) IBM parallel CTC Adapter (0.0.0602)
- 3) IBM parallel CTC Adapter (0.0.0604)

```
4) IBM Hipersocket (0.0.0700)
5) IBM Hipersocket (0.0.0701)
6) IBM Hipersocket (0.0.0702)
7) IBM OSA Express Network card (0.0.050c)
8) IBM OSA Express Network card (0.0.050d)
9) IBM OSA Express Network card (0.0.050e)
10) IBM OSA Express Network card (0.0.f401)
11) IBM OSA Express Network card (0.0.f400)
12) IBM OSA Express Network card (0.0.f402)
13) IBM IUCV
```

```
> 4
```

```
Device address for read channel [0.0.700]
[0.0.700]> 0.0.700
```

```
Device address for write channel
> 0.0.701
```

```
Device address for data channel
> 0.0.702
```

Next, choose manual configuration then enter the IP address, netmask, broadcast address, IP address of the gateway, and the searchlist of the DNS server.

Example 4.12: Network Device Name

Automatic configuration via DHCP?

```
1) Yes
2) No
```

```
> 2
```

```
Enter your IP address
> 192.168.0.20
```

```
Enter your netmask. For a normal class C network, this is usually
255.255.255.0 [255.255.255.0]
> 255.255.255.0
```

```
Enter the IP address of the gateway. Leave empty if you don't need one
> 192.168.0.1
```

```
Enter your search domains, separated by a space:
> example.com
```

4.3.2 Gigabit Ethernet with the qeth Module

Select an IBM OSA Express Network card from the list of network devices, and then 1 for ethernet. When prompted, enter the network device's read, write, and data channel numbers (for example, 0.0.0600, 0.0.0601, and 0.0.0602) and the port name, if applicable. Choose whether to enable OSI Layer 2 support.

Example 4.13: *Network Device Driver Parameters*

```
Detecting and loading network drivers
netiucv.8db02b: driver initialized
```

```
Choose the network device.
```

- 1) IBM OSA Express Network card (0.0.09a0)
- 2) IBM OSA Express Network card (0.0.09a1)
- 3) IBM OSA Express Network card (0.0.09a2)
- 4) IBM OSA Express Network card (0.0.0600)
- 5) IBM OSA Express Network card (0.0.0601)
- 6) IBM OSA Express Network card (0.0.0602)
- 7) IBM IUCV

```
> 4
```

```
Please choose the physical medium.
```

- 1) Ethernet
- 2) Token Ring

```
> 1
```

```
Enter the relative port number
```

```
> 0
```

```
Device address for read channel
```

```
[0.0.0600]> 0.0.0600
```

```
Device address for write channel
```

```
> 0.0.0601
```

```
Device address for data channel
```

```
> 0.0.0602
```

```
Portname to use
```

```
> DT70
```

```
Enable OSI Layer 2 support?
```

- 1) Yes
- 2) No

```
> 2
```

Next, deny the DHCP configuration and enter the IP address and netmask. Now enter the IP address of the gateway (if applicable), the search domain(s) and the IP address of the DNS server.

Example 4.14: *Network configuration*

Automatic configuration via DHCP?

1) Yes

2) No

> 2

Enter your IPv4 address.

Example: 192.168.5.77/24

> 192.168.0.20

Enter your netmask. For a normal class C network, this is usually

255.255.255.0

[255.255.255.0]> 255.255.255.0

Enter the IP address of the gateway. Leave empty if you don't need one

> 192.168.0.2

Enter your search domains, separated by a space:

> example.net

Enter the IP address of your name server. Leave empty or enter "+++" if you don't need one

> 192.168.0.1

4.4 The parmfile—Automating the System Configuration

The installation process can be partly automated by specifying the crucial 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. User interaction is thus limited to the actual YaST installation controlled by YaST dialogues.

The following parameters can be passed to the installation routine, which takes them as default values for installation. All IP addresses, server names, and numerical values are just examples. Replace these values with the ones needed in your installation scenario.

The number of lines in the `parmfile` is limited to 10. Specify more than one parameter on a line. Parameter names are not case-sensitive. Separate the parameters

by spaces. You may specify the parameters in any order. Always keep the `PARAMETER=value` string together in one line. 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, enter one of the following parameters at the bootprompt: `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 pauses and asks you to enter the value manually.

4.4.1 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` implies setting `Manual` to 0.

`Info=<URL>`

Specifies a location for a file from which to read additional options. 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 Section 21.1.5, “Creating the `info` File” (page 345). Since the `Info` file can typically only be accessed through the network on System `z`, you cannot use it to specify options required to setup the network, i.e. options described in Section 4.4.2, “Configuring the Network Interface” (page 71). Also other `linuxrc` specific options such as for debugging have to be specified in the `parmfile` to be effective.

TIP: Creating a File with Autoinstallation Information

At the very end of the installation of a system you can check *Clone This System for Autoyast*. This creates a ready-to-use profile as `/root/autoinst.xml` that can be used to create clones of this particular installation. To create an autoinstallation file from scratch or to edit an

existing one, use the YaST module *Autoinstallation*. For more information about AutoYaST, refer to Chapter 21, *Automated Installation* (page 337).

4.4.2 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 given in Section “Configuring a Network Connection Manually” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

Hostname=zseries.example.com

Enter the fully qualified hostname.

Domain=example.com

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

HostIP=192.168.1.2

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 officially 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 officially supported), enter the IP address of the peer:

Pointopoint=192.168.55.20

OsaInterface=<lcs|qdio> OsaMedium=<eth|tr>

For osa network devices, specify the host interface (qdio or lcs) and the physical medium (eth for ethernet or tr for token ring).

Layer2=<0|1>

For osa QDIO ethernet and hsi devices, specify whether to enable OSI Layer 2 support.

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

For Layer 2-enabled osa QDIO ethernet devices, specify the manual MAC address. Note that this is distinct from HWAddr, which contains the default MAC address as detected by linuxrc.

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 ctcc and escon (CTC and ESCON are no longer officially supported):

ReadChannel=0.0.0424
WriteChannel=0.0.0425

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

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

CTCProtocol=<0/1/2>

Valid entries would be:

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

- Network device type `osa` with interface `lcs`:

```
ReadChannel=0.0.0124
Portname=1
```

`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=PARTNER
```

Enter the name of the peer machine.

- Network device type `osa` with interface `qdio` for OSA-Express Gigabit Ethernet and OSA-Express High-speed Token Ring:

```
ReadChannel=0.0.0524
WriteChannel=0.0.0525
DataChannel=0.0.0526
Portname=FEF400
```

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

- Interface `hsi` for HiperSockets and VM guest LANs:

```
ReadChannel=0.0.0624
WriteChannel=0.0.0625
DataChannel=0.0.0626
```

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

4.4.3 Specifying the Installation Source and YaST Interface

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

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

If an ftp or smb URL is given, specify the username and password with the URL. These parameters are optional and anonymous or guest login is assumed if they are not given.

```
Install=ftp://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/
```

```
UseSSH=1 UseVNC=1 Display_IP=192.168.42.42
```

Depending on which parameter you give, a remote X server, SSH, or VNC will be used for installation. UseSSH enables SSH installation, UseVNC 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 given address. Only one of these parameters should be set at any time.

IMPORTANT: X Authentication Mechanism

The direct installation with the X Window System relies on a primitive authentication mechanism based on hostnames. This mechanism is disabled on current SUSE Linux Enterprise Server versions. Installation with 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:

```
SSHPassword=<a password>
```

4.4.4 Example Parmfiles

For an automatic installation with AutoYaST in an LPAR, it is preferable that the parmfile has just one long line. If multiple lines are desired for readability, use blank characters at the beginning and end of each line. The maximum number of lines in a parmfile is 10.

To receive potential error messages on the console, use

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

Example 4.15: *Parmfile for Installation with NFS, VNC, and IUCV and AutoYaST with HTTP*

```
ramdisk_size=131072 root=/dev/ram1 ro init=/linuxrc TERM=dumb
instnetdev=iucv iucvpeer=ROUTER01 pointopoint=192.168.0.1
hostip=192.168.0.2
nameserver=192.168.0.3
install=nfs://192.168.0.4/SLES/SLES-11-s390x/DVD1
autoyast=http://192.168.0.5/autoinst.xml
linuxrclog=/dev/console usevnc=1
vncpassword=testin
```

Example 4.16: *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=zseries.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-11-s390x/DVD1/
UseSSH=1 SSHPassword=testing linuxrclog=/dev/console
```

4.5 Using the vt220 Terminal Emulator

Recent MicroCode Levels allow the use of an integrated vt220 terminal emulator in addition to the standard line mode terminal. The vt220 terminal is connected to /dev/ttyS1. The line mode terminal is connected to /dev/ttyS0. If the vt220 emulation is available, an icon for an integrated vt220 ASCII console appears next to the icon for the 3215 console on the HMC/SE.

To activate vt220 support on your machine, edit /etc/inittab as user root. Look for the following line and delete the leading # sign:

```
#2:2345:respawn:/sbin/mingetty --noclear /dev/ttyS1 xterm
```

Save the file and run `telinit q` to pass the changes in /etc/inittab to init. The vt220 terminal should then be ready to use. If not, try hitting Enter at the terminal until the login prompt appears.

Make sure that you do not apply the changes as described above to a system that does not support vt220 terminal emulators. Otherwise, login might become impossible on this system and you will be shown the following message:

```
INIT respawning too fast, disabled for 5 minutes.
```

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=ttyS0 console=ttyS1
```

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

```
parameters = "root=/dev/dasda2 TERM=dumb console=ttyS0 console=ttyS1"
```

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

4.6 Further In-Depth Information about IBM System z

IBM has published a number of very interesting documents about their System z platform. Find them at <http://www.redbooks.ibm.com>.

4.6.1 IBM System z with SUSE Linux Enterprise Server

Find additional in-depth technical documentation about the kernel and application topics on IBM System z with SUSE Linux Enterprise Server at the following location:

- http://www.ibm.com/developerworks/linux/linux390/documentation_novell_suse.html

4.6.2 Hardware

For a first glance at the technical details of some systems, refer to:

- IBM System z10 Enterprise Class Technical Introduction (SG24-7515)

- IBM System z9 Business Class Technical Introduction (SG24-7241)
- Linux on zSeries Fibre Channel Protocol Implementation Guide (SG24-6344)

4.6.3 General Documents about Linux on IBM System z

A general coverage of Linux on IBM System 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.

4.6.4 Technical Issues of Linux on IBM System z

Refer to the following documents to get in-depth technical information about the Linux kernel and application topics. Refer to the Internet for up-to-date versions of these documents for the most recent code drop (<http://www.ibm.com/developerworks/linux/linux390/index.html>).

- 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 System z9-109 Technical Introduction (SG26-6669)
- IBM System z10 Enterprise Class Technical Guide (SG24-7516)

There also is a Redbook for Linux application development on <http://www.redbooks.ibm.com>:

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

4.6.5 Advanced Configurations for Linux on IBM System z

Refer to the following Redbooks, Redpapers, and links for some more complex IBM System 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

<http://publibz.boulder.ibm.com/epubs/pdf/hcsg4a00.pdf>

- Linux on System z documentation for "Development stream"

http://www.ibm.com/developerworks/linux/linux390/development_documentation.html

Part II. Manual Deployment

Deployment Strategies

There are several different ways to deploy SUSE Linux Enterprise Server. Choose from various approaches ranging from a local installation using physical media or a network installation server to a mass deployment using a remote-controlled, highly-customized, and automated installation technique. Select the method that best matches your requirements.

5.1 Deploying up to 10 Workstations

If your deployment of SUSE Linux Enterprise Server only involves 1 to 10 workstations, the easiest and least complex way of deploying SUSE Linux Enterprise Server is a plain manual installation as featured in Chapter 6, *Installation with YaST* (page 93). Manual installation can be done in several different ways, depending on your requirements:

Installing from the SUSE Linux Enterprise Server Media (page 82)

Consider this approach if you want to install a single, disconnected workstation.

Installing from a Network Server Using SLP (page 82)

Consider this approach if you have a single workstation or a small number of workstations and if a network installation server announced via SLP is available.

Installing from a Network Server (page 83)

Consider this approach if you have a single workstation or a small number of workstations and if a network installation server is available.

Table 5.1: *Installing from the SUSE Linux Enterprise Server Media*

Installation Source	SUSE Linux Enterprise Server media kit
Tasks Requiring Manual Interaction	<ul style="list-style-type: none">• Inserting the installation media• Booting the installation target• Changing media• Determining the YaST installation scope• Configuring the system with YaST system
Remotely Controlled Tasks	None
Details	“Installing from the SUSE Linux Enterprise Server Media (DVD, CD, USB)” (page 93)

Table 5.2: *Installing from a Network Server Using SLP*

Installation Source	Network installation server holding the SUSE Linux Enterprise Server installation media
Tasks Requiring Manual Interaction	<ul style="list-style-type: none">• Inserting the boot disk• Booting installation target• Determining the YaST installation scope• Configuring the system with YaST

Remotely Controlled Tasks	None, but this method can be combined with VNC
Details	Section 6.1.1, “Installing from a Network Server Using SLP” (page 96)

Table 5.3: *Installing from a Network Server*

Installation Source	Network installation server holding the SUSE Linux Enterprise Server installation media
Tasks Requiring Manual Interaction	<ul style="list-style-type: none"> • Inserting the boot disk • Providing boot options • Booting the installation target • Determining the YaST installation scope • Configuring the system with YaST
Remotely Controlled Tasks	None, but method can be combined with VNC
Details	Section 6.1.2, “Installing from a Network Source without SLP” (page 96)

5.2 Deploying up to 100 Workstations

With a growing number of workstations to install, you certainly do not want to install and configure each one of them manually. There are many automated or

semiautomated approaches as well as several options for performing an installation with minimal to no physical user interaction.

Before considering a fully-automated approach, take into account that the more complex the scenario gets the longer it takes to set up. If a time limit is associated with your deployment, it might be a good idea to select a less complex approach that can be carried out much more quickly. Automation makes sense for huge deployments and those that need to be carried out remotely.

Choose from the following options:

Simple Remote Installation via VNC—Static Network Configuration (page 85)

Consider this approach in a small to medium scenario with a static network setup. A network, network installation server, and VNC viewer application are required.

Simple Remote Installation via VNC—Dynamic Network Configuration (page 86)

Consider this approach in a small to medium scenario with dynamic network setup through DHCP. A network, network installation server, and VNC viewer application are required.

Remote Installation via VNC—PXE Boot and Wake on LAN (page 86)

Consider this approach in a small to medium scenario that needs to be installed via the network and without physical interaction with the installation targets. A network, a network installation server, network boot images, network bootable target hardware, and a VNC viewer application are required.

Simple Remote Installation via SSH—Static Network Configuration (page 87)

Consider this approach in a small to medium scenario with static network setup. A network, network installation server, and SSH client application are required.

Remote Installation via SSH—Dynamic Network Configuration (page 88)

Consider this approach in a small to medium scenario with dynamic network setup through DHCP. A network, network installation server, and SSH client application are required.

Remote Installation via SSH—PXE Boot and Wake on LAN (page 88)

Consider this approach in a small to medium scenario that needs to be installed via the network and without physical interaction with the installation targets. A

network, a network installation server, network boot images, network bootable target hardware, and an SSH client application are required.

Simple Mass Installation (page 89)

Consider this approach for large deployments to identical machines. If configured to use network booting, physical interaction with the target systems is not needed at all. A network, a network installation server, a remote controlling application (such as a VNC viewer or an SSH client), and an AutoYaST configuration profile are required. If using network boot, a network boot image and network bootable hardware are required, as well.

Rule-Based Autoinstallation (page 90)

Consider this approach for large deployments to various types of hardware. If configured to use network booting, physical interaction with the target systems is not needed at all. A network, a network installation server, a remote controlling application (such as a VNC viewer or an SSH client), and several AutoYaST configuration profiles as well (as a rule setup for AutoYaST) are required. If using network boot, a network boot image and network bootable hardware are required, as well.

Table 5.4: *Simple Remote Installation via VNC—Static Network Configuration*

Installation Source	Network
Preparations	<ul style="list-style-type: none">• Setting up an installation source• Booting from the installation media
Control and Monitoring	Remote: VNC
Best Suited For	Small to medium scenarios with varying hardware
Drawbacks	<ul style="list-style-type: none">• Each machine must be set up individually• Physical access is needed for booting

Details	Section 14.1.1, “Simple Remote Installation via VNC—Static Network Configuration” (page 238)
---------	--

Table 5.5: *Simple Remote Installation via VNC—Dynamic Network Configuration*

Installation Source	Network
Preparations	<ul style="list-style-type: none"> • Setting up the installation source • Booting from the installation media
Control and Monitoring	Remote: VNC
Best Suited For	Small to medium scenarios with varying hardware
Drawbacks	<ul style="list-style-type: none"> • Each machine must be set up individually • Physical access is needed for booting
Details	Section 14.1.2, “Simple Remote Installation via VNC—Dynamic Network Configuration” (page 239)

Table 5.6: *Remote Installation via VNC—PXE Boot and Wake on LAN*

Installation Source	Network
Preparations	<ul style="list-style-type: none"> • Setting up the installation source • Configuring DHCP, TFTP, PXE boot, and WOL

	<ul style="list-style-type: none"> • Booting from the network
Control and Monitoring	Remote: VNC
Best Suited For	<ul style="list-style-type: none"> • Small to medium scenarios with varying hardware • Completely remote installs; cross-site deployment
Drawbacks	Each machine must be set up manually
Details	Section 14.1.3, “Remote Installation via VNC—PXE Boot and Wake on LAN” (page 241)

Table 5.7: *Simple Remote Installation via SSH—Static Network Configuration*

Installation Source	Network
Preparations	<ul style="list-style-type: none"> • Setting up the installation source • Booting from the installation media
Control and Monitoring	Remote: SSH
Best Suited For	<ul style="list-style-type: none"> • Small to medium scenarios with varying hardware • Low bandwidth connections to target
Drawbacks	<ul style="list-style-type: none"> • Each machine must be set up individually • Physical access is needed for booting

Details	Section 14.1.4, “Simple Remote Installation via SSH—Static Network Configuration” (page 242)
---------	--

Table 5.8: *Remote Installation via SSH—Dynamic Network Configuration*

Installation Source	Network
Preparations	<ul style="list-style-type: none"> • Setting up the installation source • Booting from installation media
Control and Monitoring	Remote: SSH
Best Suited For	<ul style="list-style-type: none"> • Small to medium scenarios with varying hardware • Low bandwidth connections to target
Drawbacks	<ul style="list-style-type: none"> • Each machine must be set up individually • Physical access is needed for booting
Details	Section 14.1.5, “Simple Remote Installation via SSH—Dynamic Network Configuration” (page 243)

Table 5.9: *Remote Installation via SSH—PXE Boot and Wake on LAN*

Installation Source	Network
Preparations	<ul style="list-style-type: none"> • Setting up the installation source • Configuring DHCP, TFTP, PXE boot, and WOL

	<ul style="list-style-type: none"> • Booting from the network
Control and Monitoring	Remote: SSH
Best Suited For	<ul style="list-style-type: none"> • Small to medium scenarios with varying hardware • Completely remote installs; cross-site deployment • Low bandwidth connections to target
Drawbacks	Each machine must be set up individually
Details	Section 14.1.6, “Remote Installation via SSH—PXE Boot and Wake on LAN” (page 245)

Table 5.10: *Simple Mass Installation*

Installation Source	Preferably network
Preparations	<ul style="list-style-type: none"> • Gathering hardware information • Creating AutoYaST profile • Setting up the installation server • Distributing the profile • Setting up network boot (DHCP, TFTP, PXE, WOL) <p><i>or</i></p> <p>Booting the target from installation media</p>

Control and Monitoring	Local or remote through VNC or SSH
Best Suited For	<ul style="list-style-type: none"> • Large scenarios • Identical hardware • No access to system (network boot)
Drawbacks	Applies only to machines with identical hardware
Details	Section 21.1, “Simple Mass Installation” (page 337)

Table 5.11: *Rule-Based Autoinstallation*

Installation Source	Preferably network
Preparations	<ul style="list-style-type: none"> • Gathering hardware information • Creating AutoYaST profiles • Creating AutoYaST rules • Setting up the installation server • Distributing the profile • Setting up network boot (DHCP, TFTP, PXE, WOL) <p><i>or</i></p> <p>Booting the target from installation media</p>
Control and Monitoring	Local or remote through SSH or VNC

Best Suited For	<ul style="list-style-type: none"> • Varying hardware • Cross-site deployments
Drawbacks	Complex rule setup
Details	Section 21.2, “Rule-Based Autoinstallation” (page 349)

5.3 Deploying More than 100 Workstations

Most of the considerations brought up for medium installation scenarios in Section 5.1, “Deploying up to 10 Workstations” (page 81) still hold true for large scale deployments. However, with a growing number of installation targets, the benefits of a fully automated installation method outweigh its drawbacks.

It pays off to invest a considerable amount of time to create a sophisticated rule and class framework in AutoYaST to match the requirements of a huge deployment site. Not having to touch each target separately can save you a tremendous amount of time depending on the scope of your installation project.

As an alternative, and if user settings should be done during the first bootup, create preload images with kiwi and firstboot. Deploying such images could even be done by a PXE boot server specialized for this task. For more details, see Chapter 17, *KIWI* (page 307), Chapter 21, *Automated Installation* (page 337), and Chapter 20, *Deploying Customized Preinstallations* (page 323).

Installation with YaST

After your hardware has been prepared for the installation of SUSE® Linux Enterprise Server as described in Part I, “Architecture Specific Installation Considerations” (page 5) and after the connection with the installation system has been established, you are presented with the interface of SUSE Linux Enterprise Server’s system assistant YaST. YaST guides you through the entire installation and configuration procedure.

6.1 Choosing the Installation Method

After having selected the installation medium, determine the suitable installation method and boot option that best matches your needs:

Installing from the SUSE Linux Enterprise Server Media (DVD, CD, USB)

Choose this option if you want to perform a stand-alone installation and do not want to rely on a network to provide the installation data or the boot infrastructure. The installation proceeds exactly as outlined in Section 6.2, “The Installation Workflow” (page 96).

Installing from a Network Server

Choose this option if you have an installation server available in your network or want to use an external server as the source of your installation data. This setup can be configured to boot from physical media (Floppy, CD/DVD, or hard disk)

or configured to boot via network using PXE/BOOTP. Refer to Section 6.1.1, “Installing from a Network Server Using SLP” (page 96), Section 6.1.2, “Installing from a Network Source without SLP” (page 96), or Chapter 14, *Remote Installation* (page 237) for details.

SUSE Linux Enterprise Server supports several different boot options from which you can choose, depending on the hardware available and on the installation scenario you prefer. Booting from the SUSE Linux Enterprise Server media is the most straightforward option, but special requirements might call for special setups:

Table 6.1: *Boot Options*

Boot Option	Description
DVD	This is the easiest boot option. This option can be used if the system has a local DVD-ROM drive that is supported by Linux.
USB Mass Storage Device	<p>In case your machine is not equipped with an optical drive, you can boot the installation image from a USB mass storage device such as a USB stick. To create a bootable USB storage device, you need to copy either the DVD or the Mini CD iso image to the device using the <code>dd</code> command (the USB device must not be mounted, all data on the device will be erased):</p> <pre>dd if=PATH_TO_ISO_IMAGE of=USB_STORAGE_DEVICE bs=4M</pre> <p><code>dd</code> is available on Linux and MacOS by default. A Microsoft Windows* version can be downloaded from http://www.chrysocome.net/dd.</p>

Boot Option	Description
	<p>IMPORTANT: Compatibility</p> <p>Please note that booting from a USB Mass Storage Device is <i>not</i> supported on UEFI machines (this includes the complete ia64 architecture) and on the ppc64 architecture.</p>
PXE or BOOTP	<p>Booting over the network must be supported by the system's BIOS or firmware, and a boot server must be available in the network. This task can also be handled by another SUSE Linux Enterprise Server system. Refer to Chapter 14, <i>Remote Installation</i> (page 237) for more information.</p>
Hard Disk	<p>SUSE Linux Enterprise Server installation can also be booted from the hard disk. To do this, copy the kernel (<code>linux</code>) and the installation system (<code>initrd</code>) from the directory <code>/boot/architecture/</code> on the installation media to the hard disk and add an appropriate entry to the existing boot loader of a previous SUSE Linux Enterprise Server installation.</p>

TIP: Booting from DVD on UEFI machines

❏ **amd64 em64t:** DVD1 can be used as a boot medium for machines equipped with UEFI (Unified Extensible Firmware Interface). Refer to your vendor's documentation for specific information. If booting fails, try to enable CSM (Compatibility Support Module) in your firmware. ❏

6.1.1 Installing from a Network Server Using SLP

If your network setup supports OpenSLP and your network installation source has been configured to announce itself via SLP (described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246)), boot the system, press F4 in the boot screen and select *SLP* from the menu.

The installation program configures the network connection with DHCP and retrieves the location of the network installation source from the OpenSLP server. If the automatic DHCP network configuration fails, you are prompted to enter the appropriate parameters manually. The installation then proceeds as described below with the exception of the network configuration step that is needed prior to adding additional repositories. This step is not needed as the network is already configured and active at this point.

6.1.2 Installing from a Network Source without SLP

If your network setup does not support OpenSLP for the retrieval of network installation sources, boot the system and press F4 in the boot screen to select the desired network protocol (NFS, HTTP, FTP, or SMB/CIFS). Provide the server's address and the path to the installation media.

The installation program automatically configures the network connection with DHCP. If this configuration fails, you are prompted to enter the appropriate parameters manually. The installation retrieves the installation data from the source specified. The installation then proceeds as described below with the exception of the network configuration step needed prior to adding additional repositories. This step is not needed as the network is already configured and active at this point.

6.2 The Installation Workflow

The SUSE Linux Enterprise Server installation is split into three main parts: preparation, installation, and configuration. During the preparation phase you configure some basic parameters such as language, time, hard disk setup and

installation scope. In the non-interactive installation phase the software is installed and the system is prepared for the first boot. Upon finishing the installation the machine reboots into the newly installed system and starts the final system configuration. In this stage, network and Internet access, as well as hardware components such as printers, are set up.

6.3 IBM POWER: System Start-Up for Network Installation

For IBM POWER platforms, the system is booted (IPL, Initial Program Load) as described in Section 3.2, “Preparation” (page 23). For a network installation, SUSE Linux Enterprise Server does not show a splash screen or a boot loader command line on these systems. During the installation load the kernel manually. YaST starts with its installation screen as soon as a connection has been established to the installation system via VNC, X, or SSH. Because there is no splash screen or boot loader command line, kernel or boot parameters cannot be entered on the screen, but must be included in the kernel image using the `mkzimage_cmdline` utility.

TIP: IBM POWER: The Next Steps

To install, follow the description of the installation procedure with YaST starting from Section 6.8, “Welcome” (page 105).

6.4 IBM System z: System Start-Up for Installation

For IBM System z platforms, the system is booted (IPL, Initial Program Load) as described in Section 4.2.4, “IPLing the SUSE Linux Enterprise Server Installation System” (page 56). 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 as soon as 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 4.4, “The `parmfile`—Automating the System Configuration” (page 69)).

TIP: IBM System z: The Next Steps

To install, follow the description of the installation procedure with YaST starting from Section 6.8, “Welcome” (page 105).

6.5 System Start-Up for Installation

You can install SUSE Linux Enterprise Server from local installation sources, such as the SUSE Linux Enterprise Server CDs or DVD, or from network source of an FTP, HTTP, NFS, or SMB server. Any of these approaches requires physical access to the system to install as well as user interaction during the installation. The installation procedure is basically the same regardless of the installation source. Any exceptions are sufficiently highlighted in the following workflow description. For a description on how to perform non-interactive, automated installations, refer to Part IV, “Automated Installations” (page 335).

6.6 The Boot Screen on Machines Equipped with Traditional BIOS

The boot screen displays a number of options for the installation procedure. *Boot from Hard Disk* boots the installed system and is selected by default, because the CD is often left in the drive. 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 “F5Kernel” (page 100) for boot options that disable potentially problematic functions.

Repair Installed System

Boots into the graphical repair system. More information on repairing an installed system is available in Section “Recovering a Corrupted System” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

Rescue System

Starts a minimal Linux system without a graphical user interface. For more information, see Section “Using the Rescue System” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

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.


Firmware Test

Starts a BIOS checker that validates ACPI and other parts of your BIOS.

Memory Test

Tests your system RAM using repeated read and write cycles. Terminate the test by rebooting. For more information, see Section “Fails to Boot” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

Figure 6.1: *The Boot Screen on Machines with a Traditional BIOS*



Use the function keys indicated in the bar 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. Select *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. 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 SLP in Chapter 22, *SLP Services in the Network* (↑*Administration Guide*).

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 <http://drivers.suse.com/>. These drivers have been created via the Partner Linux Driver Program.

TIP: Using IPv6 during 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 bootprompt: `ipv6=1` (accept IPv4 and IPv6) or `ipv6only=1` (accept IPv6 only).

After starting the installation, SUSE Linux Enterprise Server loads and configures a minimal Linux system to run the installation procedure. To view the boot messages and copyright notices during this process, press `Esc`. On completion of this process, the YaST installation program starts and displays the graphical installer.

TIP: Installation without a Mouse

If the installer does not detect your mouse correctly, use `Tab` 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 the `Tab` button.

6.6.1 Providing Data to Access an SMT Server

By default, updates for SUSE Linux Enterprise Server are delivered by the Novell Customer Center. If your network provides a so called SMT 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. This information can either be entered at the boot prompt (as described here) or during the registration process as described in Section “Local Registration Server” (page 128).

regurl

URL of the SMT server. This URL has a fixed format `https://FQN/center/regsvc/`. *FQN* has to be a fully qualified hostname of the SMT server. Example:

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

regcert

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

URL

Remote location (http, https or ftp) from which the certificate can be downloaded. Example:

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

Floppy

Specifies a location on a floppy. The floppy has to be inserted at boot time, as you will not be prompted to insert it if it is missing. The value has to start with the string `floppy` followed by the path to the certificate. Example:

```
regcert=floppy/smt/smt-ca.crt
```

local path

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

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

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 either the certificate will be installed by an add-on product, or if you are using a certificate issued by an official certificate authority. Example:

```
regcert=done
```

WARNING: Beware of typing errors

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

value for `regcert` has been entered, you will be prompted for a local path to the certificate.

In case `regcert` is not specified, it will default to `http://FQN/smt.crt` with `FQN` being the name of the SMT server.

6.6.2 Configuring an Alternative Data Server for `supportconfig`

The data that `supportconfig` (see Chapter 2, *Gathering System Information for Support* (↑*Administration Guide*) for more information) gathers is sent to the Novell 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 has to be entered at the boot prompt.

`supporturl`

URL of the server. The URL has the format `http://FQN/Path/` `FQN` has to be full qualified hostname of the server, `Path` has to be replaced with the location on the server. Example:

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

6.7 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, that prevents booting malicious code by only allowing signed boot loaders to be executed. See Chapter 11, *UEFI (Unified Extensible Firmware Interface)* (↑*Administration Guide*) for more information.

The boot manager GRUB, used to boot machines with a traditional BIOS, does not support UEFI, therefore GRUB is replaced with ELILO. If Secure Boot is enabled, a GRUB2 UEFI module is used via an ELILO compatibility layer. From an administrative and user perspective, both boot manager implementations behave the same and are referred to as `ELILO` in the following.

TIP: UEFI and Secure Boot are Supported by Default

The installation routine of SUSE Linux Enterprise automatically detects if the machine is equipped with UEFI. All installation sources also support Secure Boot. If an EFI system partition already exists on dual boot machines (from a Microsoft Windows 8 installation, for example), it will automatically be detected and used. Partition tables will be written as GPT on UEFI systems.

The boot screen displays a number of 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.

Rescue System

Starts a minimal Linux system without a graphical user interface. For more information, see Section “Using the Rescue System” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

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 6.2: *The Boot Screen on Machines with UEFI*



ELILO on SUSE Linux Enterprise Server does not support a graphical boot screen or a boot prompt. In order to add additional 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 at this time). The *Installation* entry will look similar to the following:

```
setparams 'Installation'

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

Add space separated parameters to the end of the line starting with `linuxefi`. A list of parameters is available at <https://en.opensuse.org/Linuxrc>. In the following example the installation language is set to German:

```
linuxefi /boot/x86_64/loader/linux Language=de_DE
```


To boot the edited entry, press F10. If you access the machine via serial console, press Esc + 0.

6.8 Welcome

Start the installation of SUSE Linux Enterprise Server by choosing your language. Changing the language will automatically preselect a corresponding keyboard layout. Override this proposal by selecting a different keyboard layout from the drop-down menu. The language selected here is also used to assume a time zone for the system clock. This setting—along with the selection of secondary languages to install on your system—can be modified later in the *Installation Summary*, described in Section 6.14, “Installation Settings” (page 115). For information about language settings in the installed system, see Chapter 13, *Changing Language and Country Settings with YaST* (page 229).

Read the license agreement that is displayed beneath the language and keyboard selection thoroughly. Use *License Translations...* to access translations. 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.

Figure 6.3: *Welcome*



6.9 IBM System z: Hard Disk Configuration

When installing on IBM System z platforms, the language selection dialog is followed by a dialog to configure the attached hard disks. 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 Section “Installing iSCSI Target and Initiator” (Chapter 14, *Mass Storage over IP Networks: iSCSI*, ↑Storage Administration Guide).

TIP: Adding DASD or zFCP Disks at a Later Stage


Adding DASD or zFCP disks is not only possible during the installation workflow, but also when the installation proposal is shown. To add disks at that stage, click *Expert* and scroll down. The DASD and zFCP entries are shown at the very bottom.

After adding the disks, reread the partition table. Return to the installation proposal screen and choose *Partitioning* then select *Reread Partition Table*. This updates the new partition table.

6.9.1 Configuring DASD Disks

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

Figure 6.4: IBM System z: Selecting a DASD



Specify the DASDs to use for the installation by selecting the corresponding entries in the list. Click *Select or Deselect*. Activate and make the DASDs available for the installation by selecting *Perform Action > Activate*. To format the DASDs, select *Perform Action > Format* right away or use the YaST partitioner later as described in Section 15.1, “Using the YaST Partitioner” (page 275).

6.9.2 Configuring zFCP Disks

To use zFCP disks for the SUSE Linux Enterprise Server installation, select *Configure zFCP Disks* in the selection dialog. This opens a dialog with a list of the zFCP disks available on the system. 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 list. *Get WWPNS* (World Wide Port Number) and *Get LUNs* (Logical Unit Number) return lists with available WWPNS and FCP-LUNs, respectively, to choose from. 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.

6.10 Media Check

The media check dialog only appears if you install from media created from downloaded ISOs. If you install from the original media kit, the dialog is skipped.

The media check examines the integrity of a medium. To start it, select the drive that contains the installation medium and click *Start Check*. The check can take some time.

To test multiple media, wait until a result message appears in the dialog before changing the medium. If the last medium checked is not the one you started the installation with, YaST prompts for the appropriate medium before continuing with the installation.

If using ISO images (for example, for installing add-on products), click *Check ISO File* and choose the image via the file dialog.

WARNING: Failure of Media Check

If the media check fails, your medium is damaged. Do not continue the installation because installation may fail or you may lose your data. Replace the broken medium and restart the installation process.

If the media check turns out positive, click *Next* to continue the installation.

6.11 Installation Mode

After a system analysis (where YaST probes for storage devices and tries to find other installed systems on your machine) the available installation modes are displayed.

New installation

Select this option to start a new installation from scratch.


Update

Select this option to update an existing installation to a newer version. For more information about system update, see Chapter 7, *Updating SUSE Linux Enterprise* (page 135).

Repair Installed System

Choose this option to repair a damaged system that is already installed. More information is available in Section “Recovering a Corrupted System” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

Figure 6.5: *Installation Mode*



Check *Include Add-On Products from Separate Media* to include add-on products during the installation. An add-on product can include extensions, third-party products and drivers or additional software for your system.

TIP: Installing Product Patches from an SMT Server on Installation

In case your organization provides the update channel for SUSE Linux Enterprise Server via an SMT server, it is possible to specify this channel as an Add-On product by entering its HTTP address. As a consequence, the system will be installed with the most current packages without having to apply the updates at the end of the installation.

Click *Next* to proceed. If you selected to include an add-on product, proceed with Section 6.11.1, “Add-On Products” (page 111), otherwise skip the next section and advance to Section 6.12, “Clock and Time Zone” (page 112).

6.11.1 Add-On Products

Add-on products can be installed either from a local source (CD, DVD, or directory) or from a network source (HTTP, FTP, NFS, CIFS,...). When installing from a network source, you need to configure the network first (unless you are performing a network installation—in this case the existing network configuration is used). Choose *Yes, Run the Network Setup* and proceed as described in Section 6.11.1.1, “Network Setup” (page 111). If the add-on product is available locally, select *No, Skip the Network Setup*.

Click *Next* and specify the product source. Source types available are *CD*, *DVD*, *Hard Disk*, *USB Mass Storage*, a *Local Directory* or a *Local ISO Image* (if no network was configured). If the add-on product is available on removable media, the system automatically mounts the media and reads its contents. If the add-on product is available on hard disk, choose *Hard Disk* to install from an unmounted hard drive, or *Local Directory/Local ISO Image* to install from the local file system. Add-on products may be delivered as a repository or as a set of RPM files. In the latter case, check *Plain RPM Directory*. Whenever a network is available, you can choose from additional remote sources such as HTTP, SLP, FTP, etc. It is also possible to specify a URL directly.

Check *Download Repository Description Files* to download the files describing the repository now. If unchecked, they will be downloaded once the installation starts. Proceed with *Next* and insert a CD or DVD if required. Depending on the product's content it may be necessary to accept additional license agreements.

It is also possible to configure add-on products later. Using add-on products on the installed system is described in Chapter 10, *Installing Add-On Products* (page 205).

TIP: Driver Updates

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

6.11.1.1 Network Setup

When invoking the network setup, YaST scans for available network cards. If more than one network card is found, you must choose the card to configure from the list.


If an ethernet network adapter is not already connected, a warning will open. Make sure the network cable is plugged in and choose *Yes, Use It*. If your network is equipped with a DHCP server, choose *Automatic Address Setup (via DHCP)*. To manually set up the network choose *Static Address Setup* and specify *IP Address*, *Netmask*, *Default Gateway IP*, and the *DNS Server IP*.

Some networks require the use of a proxy server to access the Internet. Tick the check box *Use Proxy for Accessing the Internet* and enter the appropriate specifications. Click *Accept* to perform the network setup. The installation procedure will continue with the add-on products or repositories setup as described in Section 6.11.1, “Add-On Products” (page 111).

6.12 Clock and Time Zone

In this dialog, select your region and time zone. Both are preselected according to the selected installation language. To change the preselected values, either use the map or the drop down lists 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.

Figure 6.6: *Clock and Time Zone*



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 only 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 syncing 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 severe 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.

If a network is already configured, you can configure time synchronization with an NTP server. Click *Change* to either alter the NTP settings or to *Manually* set the time. See Chapter 23, *Time Synchronization with NTP* (↑*Administration Guide*) for

more information on configuring the NTP service. When finished, click *Accept* to continue the installation.

NOTE: Time Cannot be Changed on IBM System z

Since the operating system is not allowed to change time and date directly, the *Change* option is not available on IBM System z.

6.13 Server Base Scenario

In SUSE Linux Enterprise Server, you can choose from three base scenarios. The selected scenario affects the package selection.

Physical Machine

Choose this scenario when installing on a “real” machine or a fully virtualized guest.


Virtual Machine

Choose this scenario when installing a paravirtualized guest.

XEN Virtualization Host

Choose this scenario when installing on a machine that should serve as a XEN host.

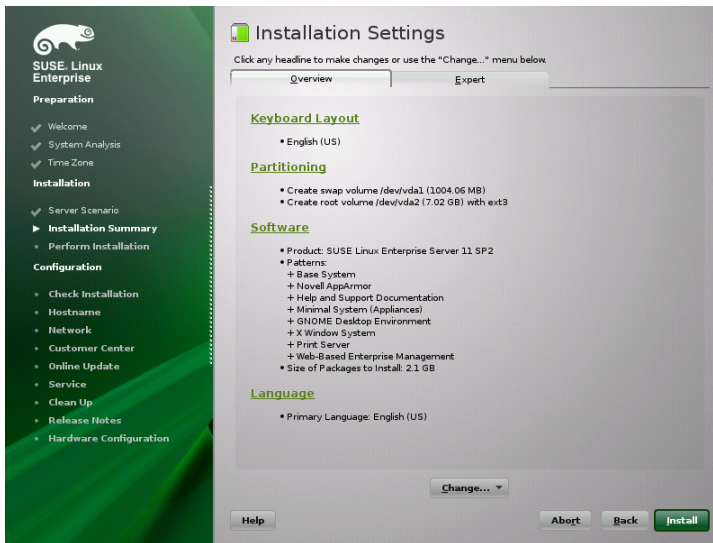
Figure 6.7: *Server Base Scenario*



6.14 Installation Settings

On the last step before the real installation takes place, you can alter installation settings suggested by YaST and also review the settings you made so far. Basic settings can be changed in the *Overview* tab, advanced options are available on the *Experts* tab. To modify the suggestions, either click *Change* and select the category to change or click on one of the headlines. After configuring any of the items presented in these dialogs, you are always returned to the Installation Settings window, which is updated accordingly.

Figure 6.8: *Installation Settings*



TIP: Restoring the Default Settings

You can reset all changes to the defaults by clicking *Change > Reset to Defaults*. YaST then shows the original proposal again.

6.14.1 Partitioning (Overview)

Review and, if necessary, change the partition setup proposed by the system. Changing the partition setup either lets you partition a specific disk or, when choosing *Custom Partitioning*, apply your own partitioning scheme. Modifying the partition setup opens the Expert Partitioner described in Section 15.1, “Using the YaST Partitioner” (page 275).

TIP: Btrfs as a default file system

The default partitioning scheme is based on Ext3 file system. To use Btrfs as a default file system instead, click *Partitioning* in the *Overview* tab and check *Use Btrfs as Default Filesystem*

NOTE: 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 but rather the ID of the physical disk. So if two or more minidisks are on the same physical disk, they all have the same ID.

To avoid problems when mounting the minidisks, always mount them either "by path" or "by UUID".

6.14.2 *Boot*ing (Expert)

▮**System z:** This module cannot be used to configure the boot loader (zipl) on the IBM System z platforms. ▮

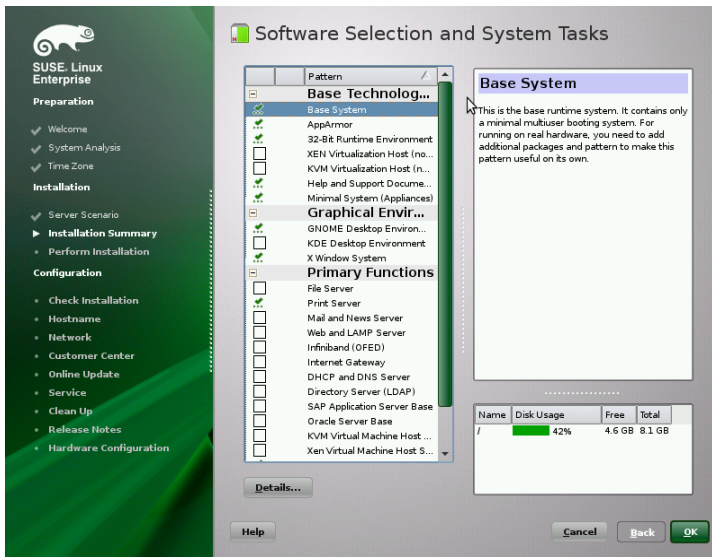
YaST proposes a boot configuration for your system. Other operating systems found on your computer, such as Microsoft Windows or other Linux installations, will automatically be detected and added to the boot loader. However, SUSE Linux Enterprise Server will be booted by default. Normally, you can leave these settings unchanged. If you need a custom setup, modify the proposal for your system. For information, see Section “Configuring the Boot Loader with YaST” (Chapter 10, *The Boot Loader GRUB*, ↑*Administration Guide*).

6.14.3 *Software* (Overview)

SUSE Linux Enterprise Server contains a number of software patterns for various application purposes. Click *Software* to start the pattern selection and modify the installation scope according to your needs. Select your pattern from the list and see a pattern description in the right part of the window. Each pattern contains a number of software packages needed for specific functions (e.g. 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.

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 Chapter 9, *Installing or Removing Software* (page 187).

Figure 6.9: *Software Selection and System Tasks*



NOTE: Default Desktop

The default desktop of SUSE Linux Enterprise Server is GNOME. To install KDE, click *Software* and select *KDE Desktop Environment* from *Graphical Environment*.

6.14.4 Language (Overview)

Here you can change the system *Language* you defined in the first step of the installation. It is also possible to add additional languages. To adjust the system language settings, select *Language*. Select a language from the list. The primary language is used as the system language. You can also adapt keyboard layout and time zone to the primary language if the current settings differ. *Details* lets you modify language settings for the user `root`, set UTF-8 support, or further specify the language (e.g. select South African English).

Choose secondary languages to be able to switch to one of these languages at any time without having to install additional packages. For more information, see Chapter 13, *Changing Language and Country Settings with YaST* (page 229).

6.14.5 *Add-On Products* (Expert)

If you added a source for an add-on media earlier, it appears here. Add, remove, or modify add-on products here, if needed. This is the same configuration dialog as discussed earlier in Section 6.11.1, “Add-On Products” (page 111).

6.14.6 *Keyboard Layout* (Expert)

To change the keyboard layout, select *Keyboard Layout*. By default, the layout corresponds to the language chosen for installation. Select the keyboard layout from the list. Use the *Test* field at the bottom of the dialog to check if you can enter special characters of that layout correctly. Options to fine-tune various settings are available under *Expert Mode*. When finished, click *Accept* to return to the installation summary.

6.14.7 *Time Zone* (Expert)

Adjust time zone and clock settings here. Provided a network is configured, you can also set up a Network Time Protocol (NTP) client that automatically synchronizes your computer with a time server. This is the same configuration as shown earlier in Section 6.12, “Clock and Time Zone” (page 112).

6.14.8 *Default Runlevel* (Expert)

SUSE Linux Enterprise Server can boot to different runlevels. Normally, there should be no need to change anything here, but if necessary set the default runlevel with this dialog. Refer to Section “Configuring System Services (Runlevel) with YaST” (Chapter 9, *Booting and Configuring a Linux System*, ↑*Administration Guide*) for more information about runlevel configuration.

6.14.9 *System* (Expert)

This dialog presents all the hardware information YaST could obtain about your computer. When called, the hardware detection routine 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 floppy. Advanced users can also change the PCI ID setup and Kernel Settings by choosing *Kernel Settings*.

6.14.10 *Kdump* (Expert)

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 Chapter 18, *kexec and kdump* (↑*System Analysis and Tuning Guide*).

6.15 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.

After the software installation has completed, the basic system is set up. Among others, “Finishing the Basic Installation” includes installing the boot manager, initializing fonts and more. Next YaST boots into the new Linux system to start the system configuration.

TIP: Existing SSH Host Keys

If you install SUSE Linux Enterprise Server on a machine with existing Linux installations, the installation routine automatically imports the SSH host key with the most recent access time from an existing installation.

6.15.1 IBM System z: IPLing the Installed System

In most cases, YaST automatically reboots into the installed system on the IBM System z platform. Known exceptions to this are installations wherein the boot

loader resides on an FCP device in environments with LPAR on a machine older than z9 or with z/VM older than release 5.3. The boot loader gets written to the device that holds the `/boot` directory. If `/boot` is not on a separate partition, it is on the same partition as the root file system `/`.

In cases where an automatic reboot is not possible, YaST will show a dialog box 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 System z HMC, select *Load*, select *Clear*, then enter the loading address (the device address of the device holding the `/boot` 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 as well as WWPN and LUN of the boot device. Now start the loading process.

z/VM Installation

Log in to the VM guest (see Example 4.5, “Configuration of a z/VM Directory” (page 55) 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.

6.15.2 IBM System z: Connecting to the Installed System

After IPLing the installed system, establish a connection with it to complete the installation. The steps involved in this vary depending on the type of connection used at the outset.

6.15.2.1 Using VNC to Connect

A message in the 3270 terminal asks you to connect to the Linux system using a VNC 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. If nothing happens for five minutes, try to initiate a connection to the Linux system using a VNC viewer.

If you connect using a Java-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/
```

6.15.2.2 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.

6.15.2.3 Using SSH to Connect

IMPORTANT: IBM System z: Connecting from a Linux or UNIX System

Start SSH in an xterm. Other terminal emulators lack complete support for the text-based interface of YaST.

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.

Once 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 may vary depending on server settings).

When the connection is established, execute the command `/usr/lib/YaST2/startup/YaST2.ssh`. Just executing the command `yast` does not suffice in this case.

YaST then begins completing the installation of the remaining packages and creating an initial system configuration.

6.16 Configuration of the Installed System

The system is now installed, but not yet configured for use. The hardware, the network and other services are not yet set up.

6.16.1 System Configuration

Having rebooted, the system starts the manual configuration. If the configuration fails at one of the steps of this stage, it restarts and continues from the last successful step.

6.16.1.1 Password for the System Administrator “root”

`root` is the name of the superuser, the administrator of the system. Unlike regular users, who may or may not have permission to execute administrative commands on the system, `root` has unlimited command capacity, for instance changing the system configuration, installing programs, and setting up new hardware. If users forget their passwords or have other problems with the system, `root` can help. The `root` account should only be used for system administration, maintenance, and repair. Logging in as `root` for daily work is rather risky: a single mistake could lead to the irretrievable loss of system files.

For verification purposes, the password for `root` must be entered twice. Do not forget the `root` password. Once entered, this password cannot be retrieved.

When typing passwords, the characters are replaced by dots, so you do not see the string you are typing. If you are unsure whether you have typed the correct string, use the *Test Keyboard Layout* field for testing purposes.

SUSE Linux Enterprise Server can use the DES, MD5, or Blowfish encryption algorithms for passwords. The default encryption type is Blowfish. To change the encryption type, click *Expert Options > Encryption Type* and select the new type.

The `root` can be changed any time later in the installed system. To do so run YaST and start *Security and Users > User and Group Management*.

6.16.1.2 Hostname and Domain Name

The hostname is the computer's name in the network. The domain name is the name of the network. A hostname and domain are proposed by default. If your system is part of a network, the hostname has to be unique in this network, whereas the domain name has to be common to all hosts on the network.

In many networks, the system receives its name over DHCP. In this case it is not necessary to modify the proposed hostname and domain name. Select *Change Hostname via DHCP* instead. To be able to access your system using this hostname, even when it is not connected to the network, select *Assign Hostname to Loopback IP*. Do not enable this option when your machine provides network services. If you often change networks without restarting the desktop environment (e.g. when switching between different WLANs), do not enable this option either, because the desktop system may get confused when the hostname in `/etc/hosts` changes.

To change hostname settings at any time after installation, use YaST *Network Devices > Network Settings*. For more information, see Section “Configuring the Network Card with YaST” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

6.16.1.3 Network Configuration

TIP: IBM System z: Network Configuration

For the IBM System z platforms, a working network connection is needed at installation time to connect to the target system, the installation source, and the YaST terminal controlling the process. The steps to set up the network are discussed in Section 4.2.5, “Network Configuration” (page 61). The IBM System z platforms only support the types of network interfaces mentioned there (OSA Ethernet, OSA Gigabit Ethernet, OSA Express Fast

Ethernet, Escon, and IUCV). The YaST dialog simply displays the interface with its settings as already configured. Just confirm this dialog to continue.

By default, *Traditional Method without NetworkManager Applet* is enabled. If desired, you can also use NetworkManager to manage all your network devices. However, the traditional method is the preferred option for server solutions. Find detailed information about NetworkManager in Chapter 26, *Using NetworkManager* (↑*Administration Guide*).

The network can also be configured after the system installation has been completed. If you skip it now, your system is left offline unable to retrieve any available updates. To configure your network connection later, select *Skip Configuration* and click *Next*.

The following network settings can be configured in this step:

General Network Settings

Enable or disable the use of NetworkManager as described above. Also change the IPv6 support here. By default the IPv6 support is enabled. To disable it, click *Disable IPv6*. For more information about IPv6, see Section “IPv6—The Next Generation Internet” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

Firewall

By default SuSEFirewall2 is enabled on all configured network interfaces. To globally disable the firewall for this computer, click on *Disable*. If the firewall is enabled, you may *Open* the SSH port in order to allow remote connections via secure shell. To open the detailed firewall configuration dialog, click on *Firewall*. See Section “Configuring the Firewall with YaST” (Chapter 15, *Masquerading and Firewalls*, ↑*Security Guide*) for detailed information.

Network Interfaces

All network cards detected by YaST are listed here. If you have already set up a network connection during the installation (as described in Section 6.11.1.1, “Network Setup” (page 111)) the card used for this connection is listed as *Configured*. A click on *Network Interfaces* opens the *Network Settings* dialog, where you can change existing configurations, set up networks cards not configured yet, or add and configure additional cards.

DSL Connections, ISDN Adapters, and Modems

If your computer is equipped with an internal DSL modem, an internal ADSL Fritz Card, an ISDN card or a modem, clicking on the respective

headline opens the configuration dialog. Refer to Chapter 11, *Accessing the Internet* (page 209) for further information.

VNC Remote Administration

To enable remote administration of your machine via VNC, click *VNC Remote Administration*. Choose *Allow Remote Administration* in the following dialog and adjust your firewall settings accordingly.

Proxy

If you have a proxy server controlling the Internet access in your network, configure the proxy URLs and authentication details in this dialog.

TIP: Resetting the Network Configuration to the Default Values

Reset the network settings to the original proposed values by clicking *Change > Reset to Defaults*. This discards any changes made.

Test Internet Connection

After having configured a network connection, you can test it. For this purpose, YaST establishes a connection to the SUSE Linux Enterprise Server server and downloads the latest release notes. Read them at the end of the installation process. A successful test is also a prerequisite for the automatic addition of the default repositories and for updating online.

If you have multiple network interfaces, verify that the desired card is used to connect to the Internet. If not, click *Change Device*.

To start the test, select *Yes, Test Connection to the Internet* and click *Next*. In the following dialog, view the progress of the test and the results. Detailed information about the test process is available via *View Logs*. If the test fails, click *Back* to return to the network configuration to correct your entries.

Proceed with *Next*. If the test was successful, the official software repositories for SUSE Linux Enterprise Server and the update repository will be configured. Downloading the repository data for the first time may take some time.

If you do not want to test the connection at this point, select *No, Skip This Test* then *Next*. This also skips downloading the release notes, configuring the customer center and updating online. These steps can be performed any time after the system has been initially configured.

6.16.1.4 Novell Customer Center Configuration

To get technical support and product updates, you need to register and activate your product with the Novell Customer Center. The *Novell Customer Center Configuration* provides assistance for doing so. Find detailed information about Novell Customer Center at <http://www.novell.com/documentation/ncc/>.

If you are offline or want to skip this step, select *Configure Later*. This also skips SUSE Linux Enterprise Server's online update.

In *Include for Convenience*, select whether to send unsolicited additional information, such as your *Hardware Profile* or *Optional Information* when registering. This simplifies the registration process. Click on *Details* to get in-depth information about how the data will be collected. In order to obtain information about which data will be sent for your specific product, the Novell server will be connected. Upon this initial connect no data other than the ID of your product will be sent to the Novell servers.

In order to become entitled for support, make sure to check *Registration Code*. You will be prompted to enter the code when proceeding with *Next*. Find more information about the technical support at http://www.novell.com/products/server/services_support.html.

NOTE: Data Privacy

No information is passed to anyone outside Novell. The data is used for statistical purposes and to enhance your convenience regarding driver support and your Web account. Find a link to the detailed privacy policy by clicking on *Details*. View the information transmitted in the log file at `root/.suse_register.log`.

Apart from activating and registering your product, this module also adds the official update repositories to your configuration. These repositories provide fixes for known bugs or security issues which can be installed via an online update.

To keep your repositories valid, select *Regularly Synchronize with Customer Center*. This option checks your repositories and adds newly available catalogs or removes obsolete ones. It does not affect manually-added repositories.

Proceed with *Next*. A connection with the Novell server is established. Follow the on-screen instructions to finish the registration.

TIP: Re-registering an Installed System with a Different Registration Code

When you register a system in Novell Customer Center, registration data is stored locally and in the Novell Customer Center database. Although it is normally not necessary, there are corner cases which may require you to re-register an already installed machine with a different registration code. To do so, proceed with the following steps on the installed system:

1. Enter the following command as user `root` to delete the installation data on the local machine:

```
suse_register.pl --erase-local-regdata
```

2. Next you need to remove the registered system from the Novell Customer Center database. Go to <http://www.suse.com/> in a browser and click *Support > Customer Center*. Log in and navigate to *My Systems > System*. Select the system and remove it by clicking on the dash sign in the bottom bar of the table.
3. Now you can re-register the machine with either `suse-register` or the YaST module *Online Update Configuration*.

Local Registration Server

If your organization provides a local registration server instead of using the Novell Customer Center, you need to specify 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. Open the dialog with *Advanced > Local Registration Server*

Registration Server

URL of the registration server. The URL has a fixed format `https://FQN/center/regsvc/` `FQN` has to be full qualified hostname of the registration server. Example:

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

Server CA certificate location

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

URL

Remote location (http, https or ftp) from which the certificate can be downloaded. Example:

```
http://smt.example.com/smt-ca.crt
```

Floppy

Specifies a location on a floppy. The floppy has to be inserted before proceeding. The value has to start with the string `floppy` followed by the path to the certificate. Example:

```
floppy/smt/smt-ca.crt
```

local path

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

```
/data/inst/smt/smt-ca.crt
```

Interactive

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

```
ask
```

Deactivate certificate installation

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

```
done
```

6.16.1.5 Online Update

If an Internet connection has been established, and updates are available, select whether to perform a YaST online update. If there are any patched packages available on the servers, download and install them now to fix known bugs or security issues. For detailed instructions see Chapter 1, *YaST Online Update* (↑*Administration Guide*). Directives on how to perform an online update in the installed system are available at Section 9.4, “Keeping the System Up-to-date” (page 199) or Chapter 1, *YaST Online Update* (↑*Administration Guide*). This step is skipped if no updates are available or no Internet connection has been established. Patches fixing security issues and recommended patches applying to your installation are automatically preselected. Click *Accept* to install them and *Next* to proceed with the system configuration.

IMPORTANT: Downloading Software Updates

The download of updates might take quite some time, depending on the bandwidth of the Internet connection and the size of the update files. In case the patch system itself is updated, the online update will restart and download more patches after the restart. If the kernel was updated, the system will reboot before completing the configuration.

6.16.1.6 Services Configuration

After testing the Internet connection and downloading the first updates, a dialog opens in which to enable and configure three network services.

CA Management

The purpose of a CA (certificate authority) is to guarantee a trust relationship among all network services communicating with each other. Without a CA, you can secure server communications with SSL and TLS separately for each individual service. By default, a CA is created and enabled during the installation. Find details about the creation of a CA with YaST in Chapter 17, *Managing X.509 Certification* (↑*Security Guide*).

OpenLDAP Server

You can run an LDAP service on your host to have a central facility manage a range of configuration files. Typically, an LDAP server handles user account data, but with SUSE Linux Enterprise Server it can also be used for mail, DHCP, and DNS data. By default, an LDAP server is set up during the installation. If you decide against the use of an LDAP server, the YaST mail server module does not work because it depends on LDAP functionality. However, you can still set up a mail server on your system with the help of the *Mail Transfer Agent* module. Find details about LDAP and its configuration with YaST in Chapter 4, *LDAP—A Directory Service* (↑*Security Guide*).

Services

The CIM (Common Information Model) Server is started by default. Click *Disable* to prevent the server automatically starting at boot time. For more information on CIM services refer to Chapter 33, *Web Based Enterprise Management Using SFCB* (↑*Administration Guide*).

If preferred, you can skip this configuration proposal for now. After the installation is finished, configure and start the same services with the help of YaST.

TIP: Resetting the Service Configuration to Defaults

Restore the defaults by clicking *Change > Reset to Defaults*. This discards any changes made.

6.16.1.7 User Authentication Method

If network access was configured successfully during the previous steps of the installation, you can now choose from several user management options. If a network connection has not been configured, create local user accounts. You may also, if present, import users from a previous installation. Also change the password encryption type in this dialog.

You can also add additional user accounts or change the user authentication method in the installed system. For detailed information about user management, see Chapter 12, *Managing Users with YaST* (page 213).

The default authentication method is *Local (/etc/passwd)*. If a former version of SUSE Linux Enterprise Server or another system using */etc/passwd* is detected, you may import local users. To do so, check *Read User Data from a Previous Installation* and click *Choose*. In the next dialog, select the users to import and finish with *OK*.

Manually enter local users by clicking *Next*. The New Local User dialog opens. After entering the first name and last name, either accept the proposal or specify a new *Username* that will be used to log in. Finally, enter a password for the user. Reenter it for confirmation (to ensure that you did not type something else by mistake). To provide effective security, a password should be between five and eight characters long. The maximum length for a password is 72 characters. However, if no special security modules are loaded, only the first eight characters are used to discern the password. Passwords are case-sensitive. Special characters (7-bit ASCII) and the digits 0 to 9 are allowed. Other special characters like 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.

Two additional options are available:

Receive System Mail

Checking this box sends messages created by the system services to the user. These are usually only sent to `root`, the system administrator. This option is useful for the most frequently used account, because it is highly recommended to log in as `root` only in special cases.

The mails sent by system services are stored in the local mailbox `/var/spool/mail/username`, where *username* is the login name of the selected user. To read e-mails after installation, you can use any e-mail client, for example KMail or Evolution.

Automatic Login

This option automatically logs the current user in to the system on start-up. 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 automatic login enabled, the system boots straight to your desktop with no authentication at all. If you store sensitive data on your system, you should not enable this option if the computer can also be accessed by others.

Enter more users by calling the *User Management* module described in Chapter 12, *Managing Users with YaST* (page 213).

When using a network server for user authentication, access to the following services can be configured:

LDAP

Users are administered centrally on an LDAP server for all systems in the network. More information is available in Section “Configuring an LDAP Client with YaST” (Chapter 4, *LDAP—A Directory Service*, ↑*Security Guide*).

NIS

Users are administered centrally on a NIS server for all systems in the network. See Section “Configuring NIS Clients” (Chapter 3, *Using NIS*, ↑*Security Guide*) for more information.

Windows Domain

SMB authentication is often used in mixed Linux and Windows networks. Detailed information is available in Section “Samba Server in the Network

with Active Directory” (Chapter 27, *Samba*, ↑*Administration Guide*) and Section “Configuring a Linux Client for Active Directory” (Chapter 5, *Active Directory Support*, ↑*Security Guide*).

Along with user administration via *LDAP* and *NIS*, you can use Kerberos authentication. To use it, select *Set Up Kerberos Authentication*. For more information on Kerberos, refer to Chapter 6, *Network Authentication with Kerberos* (↑*Security Guide*).

6.16.1.8 Release Notes

After completing the user authentication setup, YaST displays the release notes. Reading them is recommended, because they contain important up-to-date information which was not available when the manuals were printed. If you successfully tested the Internet connection, read the most recent version of the release notes, as fetched from SUSE Linux Enterprise Server's servers. Use *Miscellaneous > Release Notes* in YaST or start the SUSE Help Center to view the release notes after installation.

6.16.1.9 Hardware Configuration

At the end of the installation, YaST opens a dialog for the configuration of *Graphics Cards Printer* and *Sound*. Click the individual components to start the hardware configuration. For the most part, YaST detects and configures the devices automatically.

TIP: IBM System z: Hardware Configuration

On the IBM System z, there is no display that would be supported by XFree. Accordingly, you do not find a *Graphics Cards* entry on these systems.

You can skip any peripheral devices and configure them later, as described in Chapter 8, *Setting Up Hardware Components with YaST* (page 171). To skip the configuration, select *Skip Configuration* and click *Next*.

However, when setting up a desktop system you should configure the graphics card right away. Although the display settings as configured by YaST should be generally acceptable, most users have very strong preferences as far as resolution,

color depth, and other graphics features are concerned. To change these settings, select the respective item and set the values as desired.

TIP: Resetting Hardware Configuration to the Default Values

You can cancel any changes to the hardware configuration by clicking *Change > Reset to Defaults*. YaST then shows the original proposal again.

6.16.1.10 Installation Completed

After a successful installation, YaST shows the Installation Completed dialog. In this dialog, select whether to clone your newly installed system for AutoYaST. To clone your system, select *Clone This System for AutoYaST*. The profile of the current system is stored in `/root/autoyast.xml`. Cloning is selected by default.

AutoYaST is a system for installing one or more SUSE Linux Enterprise Server systems automatically without user intervention. AutoYaST installations are performed using a control file with installation and configuration data. For detailed information, refer to Chapter 21, *Automated Installation* (page 337). Finish the installation of SUSE Linux Enterprise Server with *Finish* in the final dialog.

6.17 Graphical Login

TIP: IBM System z: No Graphical Login

The graphical login is not available on IBM System z platforms.

SUSE Linux Enterprise Server is now fully installed and configured. Unless you enabled the automatic login function or customized the default runlevel, you should see the graphical login on your screen in which to enter a username and password to log in to the system. On single user systems with automatic login enabled, the desktop starts automatically.

Updating SUSE Linux Enterprise

SUSE® Linux Enterprise (SLE) provides the option of updating an existing system to the new version without completely reinstalling it. No new installation is needed. Existing data, such as home directories and system configuration, is kept intact. During the life-cycle of the product, you can apply Service Packs to increase system security, correct software defects and get access to new features. Install from a local CD or DVD drive or from a central network installation source.

7.1 Terminology

This chapter uses several terms. In order to understand the information, read the definitions below:

Backporting

Backporting is the act of adapting specific changes from a newer version of software and applying it to an older version. The most commonly used case is fixing security holes in older software components. Usually it is also part of a maintenance model to supply enhancements or (less commonly) new features.

Deltarpm

A deltarpm consists only of the binary diff between two defined versions of a package, and therefore has the smallest download size. Before being installed, the full RPM package is rebuilt on the local machine.

Downstream

A metaphor of how software is developed in the open source world (compare it with *upstream*). The term *downstream* refers to people or organisations like SUSE who integrate the source code from upstream with other software to build a distribution which is then used by end users. Thus, the software flows downstream from its developers via the integrators to the end users.

Online Migration

Updating to a Service Pack (SP) by using the online update tools (rather than the installation media) to install the respective patches. It updates all packages of the installed system to the latest state—including updates—of SP3 plus SP2 updates.

Package

A package is a compressed file in `rpm` format that contains all files for a particular program, including optional components like configuration, examples, and documentation.

Patch

A patch consists of one or more packages and may be applied by means of `deltarpm`s. It may also introduce dependencies to packages that are not installed yet.

Service Packs (SP)

Combines several patches into a form which is easy to install or deploy. Service packs are numbered and usually contain security fixes, updates, upgrades, or enhancements of programs.

Upstream

A metaphor of how software is developed in the open source world (compare it with *downstream*). The term *upstream* refers to the original project, author or maintainer of a software that is distributed as source code. Feedback, patches, feature enhancements, or other improvements flow from end users or contributors to upstream developers. They decide if the request will be integrated or rejected.

If the project members decide to integrate the request, it will show up in newer versions of the software. An accepted request will benefit all parties involved.

If a request is not accepted, it may be for different reasons. Either it is in a state which is not compliant with the project's guidelines, it is invalid, it is already integrated, or it is not in the interest or roadmap of the project. An unaccepted request makes it harder for upstream developers as they have to keep their

patches in sync with the upstream code. This practice is generally avoided, but sometimes it is still needed.

Update

Installation of a newer *minor* version of a package.

Upgrade

Installation of a newer *major* version of a package or distribution, which brings *new features*.

7.2 The SUSE Linux Enterprise 11 Maintenance Model

The SUSE Linux Enterprise 11 Maintenance Model combines flexibility and control of your service packs. It offers the following benefits:

- Makes service packs more lightweight and easier to test and deploy.
- Allows staying with older versions, but with support for the full system.
- Answers market needs in between service packs by selective enhancements and allows more updates in the general update repository. By selecting the enhancements, it mitigates longer periods between service packs.

7.2.1 Background Information

Over the last few years, with a clear desire for improvements based on customer feedback, SUSE has implemented various changes regarding the way we deliver updates to our users:

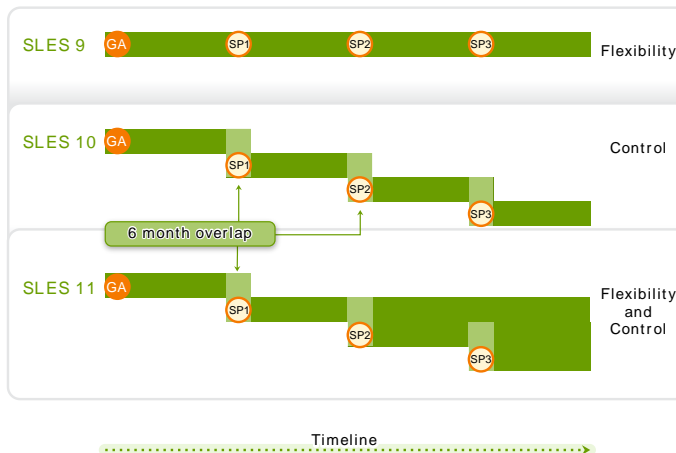
- In SLES 9, there was only one update repository that collected all the updates; the most recent release update was the only one to be supported.
- Starting with SLES 10 SP1, the concept of a “SP-specific repository” was introduced. This means that all the updates for a specific service pack are delivered in one specific repository. Once users migrate to a newer service pack, they lose access to the preceding repositories if they registered directly at the Novell Customer Center. Users of SMT or SUSE Manager were able and are still able to subscribe to any SP channel freely. The main reason for this change was the

concept of a 6-month overlap period (n-1 service pack support) to allow validation of the released service pack and a window of migration for customers, whereby they would continued to be maintained and supported fully within the old SP.

- SLES 11 GA and SLES 11 SP1 followed the SLES 10 model. With SLES 11 SP2 we introduced a new repository model, consisting of the following:
 - i. SLES 11 SP1 Updates repository remained subscribed. All updates that were also applicable to SP2 were also or only released into the SP1 Updates repository. This means that all the updates that don't break the ABI and API compatibility continued to be delivered here.
 - ii. SLES 11 SP2 Updates repository includes only the latest and any innovative updates that can't be delivered to the SP1 Updates repository (for various reasons). In addition to this, we introduced a core repository, which provided a “gap” for packages that were neither released into the SP1 nor the SP2 Updates repository.

Figure 7.1, “Maintenance Delivery Evolution (also applies to SLED)” (page 138) depicts some of the mentioned aspects above.

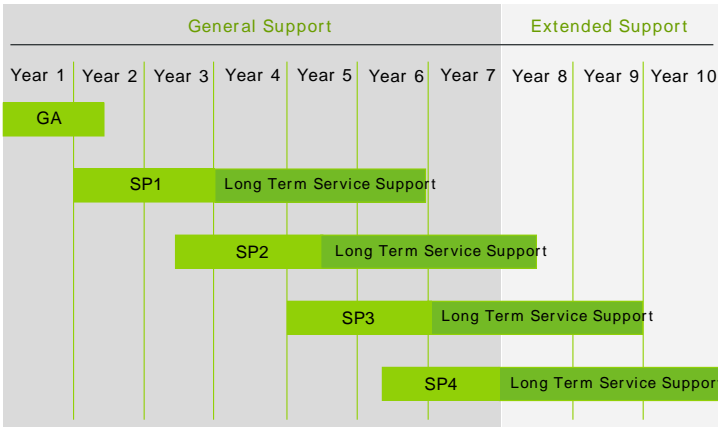
Figure 7.1: *Maintenance Delivery Evolution (also applies to SLED)*



Our products have a 10-year life-cycle: 7 years general support and 3 years extended support. Major releases are made every 4 years and service packs every 18 months.

The Long Term Service Pack Support is an extended window or extended major release life-cycle (see Figure 7.2, “Long Term Service Pack Support” (page 139)).

Figure 7.2: Long Term Service Pack Support



The Long Term Service Pack Support requires active subscription, either standard or priority. It does not affect L1 or L2 subscription terms. Security updates are handled “proactively”: these are any non-user driven critical vulnerabilities, local root exploits in Kernel or other root exploits directly executable without user interaction.

7.2.2 Support Levels

The range for extended support levels starts from year 8 and ends in year 10. These contain continued L3 engineering level diagnosis and reactive critical bug fixes. These support levels proactively update trivial local root exploits in Kernel or other root exploits directly executable without user interaction. Furthermore they support existing workloads, software stacks, and hardware with limited package exclusion list. Find an overview in Table 7.1, “Security Updates and Bug Fixes” (page 140).

Table 7.1: *Security Updates and Bug Fixes*

	— General Support —				Extended Support
Topic	Current SP	SP (n-1) 6 months	SP (n-1) with LTSS	Year 6 & 7 with LTSS	Year 8, 9, 10 with LTSS
L1/L2 Technical Services	✓	✓	✓	✓	✓
Proactive Maintenance	✓	✓		✓	
Driver Updates via PLDP	✓	✓	✓		
Proactive Security Updates	✓	✓	✓	✓	
L3 Engineering Support	✓	✓	✓	✓	✓
Back-ports available	✓		✓	✓	✓

7.2.3 Channel Model

With the former maintenance model, SUSE Linux Enterprise Server, had two channels assigned: `SLES11-SPx-Pool` and `SLES11-SPx-Updates`. During

an online migration to SPx+1 these channels were temporary replaced by SLES11-SPx-Online.

With SUSE Linux Enterprise SP 2 the channel layout has changed to support the benefits of the new maintenance model. Table 7.2, “Channel Layout for SUSE Linux Enterprise 11 SP1, SP2, and SP3” (page 141) contains a list of all channels from SP1 to SP3.

Table 7.2: *Channel Layout for SUSE Linux Enterprise 11 SP1, SP2, and SP3*

Type	SLES	SLED
Required Channels	SP1	SP1
	SLES11-SP1-Pool SLES11-SP1-Updates	SLED11-SP1-Pool SLED11-SP1-Updates
	SP2	SP2
	SLES11-SP1-Pool SLES11-SP1-Updates SLES11-SP2-Core SLES11-SP2-Updates	SLED11-SP1-Pool SLED11-SP1-Updates SLED11-SP2-Core SLED11-SP2-Updates
Optional Channels	SP3	SP3
	SLES11-SP3-Pool SLES11-SP3-Updates	SLED11-SP3-Pool SLED11-SP3-Updates
	SP1	SP1
	SLES11-SP1- Debuginfo-Pool SLES11-SP1- Debuginfo-Updates	SLED11-SP1- Debuginfo-Pool SLED11-SP1- Debuginfo-Updates
	SP2	SP2
	SLES11-SP2- Debuginfo-Core SLES11-SP2- Debuginfo-Updates	SLED11-SP2- Debuginfo-Core SLED11-SP2- Debuginfo-Updates

Type	SLES	SLED
	SLES11-Extras SLES11-SP2- Extension-Store SP3 SLES11-SP3- Debuginfo-Core SLES11-SP3- Debuginfo-Updates SLES11-SP3- Extension-Store SLES11-Extra	SLED11-Extras SLED11-SP2- Extension-Store SP3 SLED11-SP3- Debuginfo-Core SLED11-SP3- Debuginfo-Updates SLED11-SP3- Extension-Store SLED11-Extra
Product-Specific (Examples)	SLES11-WebYaST-SP2- Pool SLES11-WebYaST-SP2- Updates	SLED11-MSI-Updates

Description of Required Channels

Core

A subset of the unpacked installation media, it only contains those packages that are considered to be the “core” of SPx (approximately 30% of the package total). The SP repositories only contain packages specific to a SP and its themes (for example, hardware enablement). Exists only in SP2.

Updates

Maintenance updates to packages in the corresponding `Core` or `Pool` repository.

Pool

Containing all binary RPMs from the installation media, plus pattern information and support status metadata.

Description of Optional Channels

Debuginfo-Pool , Debuginfo-Updates

These channels contain static content. However, only the `Debuginfo-Updates` channel receives updates. Enable these channels if you need to install libraries with debug information in case of an issue.

Extension-Store

Not in use, yet. Supposed to contain packages for (future) add-on products.

LTSS-Updates

Maintenance updates to packages in the corresponding `Pool` repository for installations with Long Term Support Service (LTSS). This specific channels require an LTSS contract.

7.2.3.1 Origin of Packages

SUSE Linux Enterprise 11 SP1 The base channel for SUSE Linux Enterprise 11 SP1 is `SLES11-SP1-Pool`. This channel and its content will not change throughout the entire life cycle of SUSE Linux Enterprise 11. A package will be served from this channel as long as no update for this package is available.

Once a package gets updated, the update is served from `SLES11-SP1-Updates`. This channel contains all packages having been updated since the release of SP1.

SUSE Linux Enterprise 11 SP2 Once Service Pack 2 is installed, two additional channels will be added: `SLES11-SP2-Core` and `SLES11-SP2-Updates`. `SP2-Core` contains a subset of packages from `SP1-Pool`. These packages are newer than the versions in both SP1 channels. Once an `SP2-Core` package gets updated, the update is served from `SLES11-SP2-Updates`. This channel contains all packages having been updated since the release of SP2.

SUSE Linux Enterprise 11 SP3 With the installation of SP3 there are only two channels available: `SLES11-SP3-Pool` and `SLES11-SP3-Updates`. Any previous channels from SP2 are visible, but not enabled. These disabled channels are only needed for users who have particular needs.

7.2.3.2 Working with Channels

On registration, the system receives channels from the Novell Customer Center. The channel names map to specific URIs in the customer center (see <http://>

www.novell.com/ncc). To list all available channels on your system, use `zypper` as follows:

```
zypper repos -u
```

This gives you a list of all available channels on your system. Each channel is listed by its alias, name and whether it is enabled and will be refreshed. The option `-u` gives you also the URI from where it originated.

If you want to remove old channels (for example, from SP1), use `zypper removerepo` and the name(s) of the channel(s). For example, to remove the old SP1 and SP2 channels, use the following command:

```
zypper removerepo SLES11-SP1-Pool SLES11-SP1-Updates \  
SLE11-SP1-Debuginfo-Pool SLE11-SP1-Debuginfo-Updates \  
SLES11-SP2-Core SLES11-SP2-Updates \  
SLE11-SP2-Debuginfo-Core SLES11-SP2-Extension-Store\  
SLE11-SP2-Debuginfo-Updates
```

If you want to re-add some of your channels, log in to <http://www.novell.com/ncc> and select from the menu *My Products > Mirror Credentials*. There you can see a list of URIs; Only channels from this list of products can be added. For example, to add the SP2 Extension Store, use the following command (one line, without the backslash):

```
zypper addrepo -n SLES11-SP2-Extension-Store \  
https://nu.novell.com/repo/$RCE/SLES11-SP2-Extension-Store/  
nu_novell_com:SLES11-SP2-Extension-Store
```

7.3 Supported Upgrade Paths to SLE SP3

Upgrading SLES 8, SLES 9 and NLD 9

There is no supported direct upgrade path from these versions. Instead it is recommended to perform a new installation.

Upgrading from SUSE Linux Enterprise 10 (any Service Pack)

There is no supported direct migration path to SUSE Linux Enterprise 11 SP3. You need to do a fresh installation of SUSE Linux Enterprise 11 SP3. Proceed with Chapter 6, *Installation with YaST* (page 93).

Upgrading from SUSE Linux Enterprise 11 GA

There is no supported direct migration path to SUSE Linux Enterprise 11 SP3. An update has to be performed from SUSE Linux Enterprise 11 GA to SP1 first.

Proceed with Section 7.5, “Updating SLE 11 SP1 to SLE 11 SP2” (page 146) afterwards.

Upgrading from SUSE Linux Enterprise 11 SP1

Refer to Section 7.5, “Updating SLE 11 SP1 to SLE 11 SP2” (page 146) for details.

Upgrading from SUSE Linux Enterprise 11 SP2

Refer to Section 7.6, “Updating SLE 11 SP2 to SLE 11 SP3” (page 155) for details.

7.4 General Preparations for Updating

Before starting the update procedure, make sure your system is properly prepared. Among others, preparation involves backing up data and checking the release notes.

7.4.1 Make a Backup

Before updating, copy existing configuration files to a separate medium (such as tape device, removable hard disk, etc.) to back up the data. This primarily applies to files stored in `/etc` as well as some of the directories and files in `/var` and `/opt`. You may also want to write the user data in `/home` (the `HOME` directories) to a backup medium. Back up this data as `root`. Only `root` has read permissions for all local files.

If you have selected *Update an Existing System* as the installation mode in YaST, you can choose to do a (system) backup at a later point in time. You can choose to include all modified files and files from the `/etc/sysconfig` directory. However, this is not a complete backup, as all the other important directories mentioned above are missing. Find the backup in the `/var/adm/backup` directory.

7.4.2 Partitioning and Disk Space

Before starting your update, make note of the root partition. The command `df /` lists the device name of the root partition. In Example 7.1, “List with `df -h`” (page 146), the root partition to write down is `/dev/sda3` (mounted as `/`).

Example 7.1: *List with `df -h`*

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda3	74G	22G	53G	29%	/
tmpfs	506M	0	506M	0%	/dev/shm
/dev/sda5	116G	5.8G	111G	5%	/home
/dev/sda1	44G	4G	40G	9%	/data

Software tends to “grow” from version to version. Therefore, take a look at the available partition space with `df` before updating. If you suspect you are running short of disk space, secure your data before updating and repartitioning your system. There is no general rule regarding how much space each partition should have. Space requirements depend on your particular partitioning profile and the software selected.

7.4.3 Shut Down Virtual Machines

If your machine serves as a VM Host Server for KVM or Xen, make sure to properly shut down all running VM Guests prior to the update. Otherwise you may not be able to access the guests after the update.

7.4.4 Version Specific Requirements

For version specific requirements, refer to the release notes coming with the update product. In the release notes you can find additional information about upgrade procedures.

The current version of the release notes document containing the latest information on SUSE Linux Enterprise Server can be read online at <http://www.suse.com/doc/sles11/#additional>.

7.5 Updating SLE 11 SP1 to SLE 11 SP2

There are different supported ways for updating a SUSE Linux Enterprise 11 SP1 system to a Service Pack 2. You may either update by using the online update tools to install the respective patches (“Online Migration”) or update via the Service Pack installation media. Furthermore, updates can be performed via servers hosting Subscription Management Tool or SUSE Manager.

An Online Migration is supported by the following tools:

- *YaST wagon* (graphical user interface)
- `zypper` (command line)

Alternatively, the full Service Pack media (DVD ISO image) can be downloaded. Start the update process by booting from the physical Service Pack media or a network installation source.

7.5.1 Online Migration

Updating your system via online migration is done from within the running system. You only need to reboot once, after the update is completed.

7.5.1.1 Requirements

In order to do an online update, the following requirements must be met. Make sure to also read Section 7.4, “General Preparations for Updating” (page 145).

Product Registration

In order to be able to connect to the update channels, your product has to be registered. If this is not the case, either run the *Novell Customer Center Configuration* module in YaST or the `suse_register` command line tool to start the registration.

Run an Online Update

Make sure the currently installed version has the latest patches installed. Run an Online Update prior to the Online Migration. When using a graphical interface, start the YaST Online Update or the updater applet. On the command line, run the following commands (the last command needs to be run twice):

```
zypper ref -s
zypper update -t patch
zypper update -t patch
```

Reboot the system if needed.

See Chapter 1, *YaST Online Update* (↑*Administration Guide*) or at Section “Updating Software with Zypper” (Chapter 6, *Managing Software with*

Command Line Tools, ↑*Administration Guide*). for more information. on the online update tools.

Third-Party Software

If your setup involves third-party software or add-on software, test this procedure on another machine to make sure that the dependencies are not broken by the update.

IMPORTANT: Always Run a Complete Online Migration

The online migration always has to be completed from beginning to end. If an online migration is interrupted in between, the system is corrupted beyond recovery.

7.5.1.2 Online Migration with YaST *Wagon*

- 1 When all requirements are met (see Section 7.5.1.1, “Requirements” (page 147)), the update applet in the tray will display a message that a distribution upgrade is available. Click it to start YaST *Wagon*. Alternatively run `/usr/sbin/wagon` as root from the command line.
- 2 Confirm the *Welcome* dialog with *Next*.
- 3 If *Wagon* finds that the requirements are not met (required maintenance updates are available but not yet installed) it will do an automatic self-update, which may require a reboot. Follow the on-screen instructions.
- 4 Choose the update method in the following dialog. Choose *Customer Center* to use the default setup (recommended).

Click *Custom URL* to manually choose the software channels used for the online migration. A list of channels will be displayed, providing the opportunity to manually enable, disable, add, or delete channels. Add the SP2 update source(s). This can either be the SP2 installation media or the SP2-Core and SP2-Updates channels. Click *OK* to return to the *Update Method* dialog.

If you want to review changes to the channel setup caused by the update process, select *Check Automatic Repository Changes*.

Proceed with *Next*.

- 5 The system will be re-registered. During this process the `SP2-Core` and `SP2-Updates` channels will be added to the system (see Section 7.2.3, “Channel Model” (page 140) for more information). Confirm the addition of the channels.
- 6 If you have selected *Check Automatic Repository Changes* in the *Update Method* dialog, the list of repositories will be displayed, providing the opportunity to manually enable, disable, add, or delete channels. Proceed with *OK* when finished.
- 7 Choose the migration type:

Full migration

Updates all packages to the latest SP2 level.

Minimal Migration

Updates a minimal set of packages to the latest SP2 level.

Click *Advanced* to manually select the repositories used for upgrading.

Confirm your selection.

- 8 The *Distribution Upgrade Settings* screen opens, presenting a summary of the update configuration. The following sections are available:

Add-On Products

You may add SUSE Linux Enterprise Server add-on products or third-party products here.

Update Options

Lists the actions that will be performed during the update. You can choose whether to download all packages before installing them (default, recommended), or whether to download and install packages one by one.

Packages

Statistical overview of the update.

Backup

Set backup options.

Click *Next* and *Start The Update* to proceed.

IMPORTANT: Aborting the Online Migration

It is safe to abort the online migration on this screen *prior* to clicking *Start The Update* and on all previous screens. Click *Abort* to leave the update procedure and restore the system to the state it was in prior to starting YaST wagon. Follow the instructions on screen and perform a re-registration before leaving Wagon to remove the SP2 channels from your system.

9 During the update procedure the following steps are executed:

9a Packages will be updated.

9b `SuSEconfig` will be run.

9c The system will be rebooted (press *OK*).

9d The newly updated system will be re-registered.

10 Your system has been successfully updated to Service Pack 2.

7.5.1.3 Online Migration with zypper

1 When all requirements are met (see Section 7.5.1.1, “Requirements” (page 147)), the “products” needed for the online migration have been added to `/etc/products.d`. Get a list of these products by running the following command:

```
zypper se -t product | grep -h -- "-migration" | cut -d'|' -f2
```

This command should at least return `SUSE_SLES-SP2-migration`. Depending on the scope of your installation, more products may be listed.

2 Install the migration products retrieved on the previous step with the command `zypper in -t product LIST_OF_PRODUCTS`, for example

```
zypper in -t product SUSE_SLES-SP2-migration
```

- 3** Register the products installed in the previous step in order to get the respective update channels:

```
suse_register -d 2 -L /root/.suse_register.log'
```

- 4** Refresh repositories and services again:

```
zypper ref -s
```

- 5** Check the list of repositories you can retrieve with `zypper lr`. *At least* the following repositories need to be *Enabled*:

- SLES11-SP1-Pool
- SLES11-SP1-Updates
- SLES11-SP2-Core
- SLES11-SP2-Updates

Depending on the scope of your installation, further repositories for add-on products or extensions need to be enabled.

If one of these repositories is not enabled (the SP2 ones are not enabled by default when following this workflow), enable them with `zypper modifyrepo --enable REPOSITORY ALIAS`, for example:

```
zypper modifyrepo --enable SLES11-SP2-Core SLES11-SP2-Updates
```

If your setup contains third-party repositories that may not be compatible with SP2, disable them with `zypper modifyrepo --disable REPOSITORY ALIAS`.

- 6** Now everything is in place to perform the distribution upgrade with `zypper dup --from REPO 1 --from REPO 2 . . .`. Make sure to list all needed repositories with `--from`, for example:

```
zypper dup --from SLES11-SP2-Core --from SLES11-SP2-Updates
```

Confirm with `y` to start the upgrade.

- 7** Upon completion of the distribution upgrade from the previous step, a Minimal Migration has been performed (a minimal set of packages has been updated to the latest SP2 level). Skip this step if you do not intend to do a Full Migration.

In order to do a Full Migration (updates all packages to the latest SP2 level), run the following command:

```
zypper update -t patch
```

- 8 Now that the upgrade to SP2 has been completed, you need to re-register your product:

```
suse_register -d 2 -L /root/.suse_register.log
```

- 9 Last, reboot your system.

- 10 Your system has been successfully updated to Service Pack 2.

7.5.2 Updating by Booting from an Installation Source

As an alternative to the Online Migration (see Section 7.5.1, “Online Migration” (page 147) for details) you may also update your system by booting from an installation source—either a DVD or a network installation source. The update will start like a normal installation.

Service Pack 2 ISO images can be obtained from <http://download.novell.com/>. Either burn it to a DVD or prepare a network installation source as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246).

7.5.2.1 Updating from a Local DVD Drive

Before starting a new installation of a SUSE Linux Enterprise SP, ensure that all the Service Pack installation media (DVDs) are available.

Procedure 7.1: *Booting from the Service Pack Medium*

- 1 Insert the first SUSE Linux Enterprise SP medium and boot your machine. A boot screen similar to the original installation of SUSE Linux Enterprise 11 is displayed.
- 2 Select *Installation* and continue as outlined in the YaST installation instructions in Chapter 6, *Installation with YaST* (page 93).

7.5.2.2 Updating from a Network Installation Source

Before starting an update of a SUSE Linux Enterprise SP from a network installation source, make sure that the following requirements are met:

- A network installation source is setup according to Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246).
- A working network connection on both the installation server and the target machine that includes a name service, DHCP (optional, but needed for PXE boot), and OpenSLP (optional) exists.
- The SUSE Linux Enterprise SP DVD 1 to boot the target system *or* a target system set up for PXE boot according to Section 14.3.5, “Preparing the Target System for PXE Boot” (page 265) exist.

Please refer to Chapter 14, *Remote Installation* (page 237) for in-depth information on starting the upgrade from a remote server.

Network Installation—Boot from DVD

To perform a network installation using the SP DVD as the boot medium, proceed as follows:

- 1** Insert the SUSE Linux Enterprise SP DVD 1 and boot your machine. A boot screen similar to the original installation of SUSE Linux Enterprise 11 is displayed.
- 2** Select *Installation* to boot the SP Kernel then use **F4** to select the type of network installation source (FTP, HTTP, NFS, or SMB).
- 3** Provide the appropriate path information or select *SLP* as the installation source.
- 4** Select the appropriate installation server from those offered or use the boot options prompt to provide the type of installation source and its actual location as in Section 6.1.2, “Installing from a Network Source without SLP” (page 96). YaST starts.

Finish the installation as outlined in Section 7.5.2.3, “The Update Procedure” (page 154).

Network Installation—PXE Boot

To perform a network installation of a SUSE Linux Enterprise Service Pack via network, proceed as follows:

- 1 Adjust the setup of your DHCP server to provide the address information needed for PXE boot according to Section 14.3.5, “Preparing the Target System for PXE Boot” (page 265).

- 2 Set up a TFTP server to hold the boot image needed for PXE boot.

Use the first CD or DVD of your SUSE Linux Enterprise Service Pack for this or follow the instructions in Section 14.3.2, “Setting Up a TFTP Server” (page 258).

- 3 Prepare PXE boot and Wake-on-LAN on the target machine.
- 4 Initiate the boot of the target system and use VNC to remotely connect to the installation routine running on this machine. See Section 14.5.1, “VNC Installation” (page 270) for more information.
- 5 Finish the installation as outlined in Section 7.5.2.3, “The Update Procedure” (page 154).

7.5.2.3 The Update Procedure

Once you have successfully booted from installation medium or the network, proceed as follows to start the update:

- 1 On the *Welcome* screen choose *Language* and *Keyboard* and Accept the license agreement. Proceed with *Next*.
- 2 In case you have booted from a physical medium, perform a *Media Check* to verify the integrity of the medium. Only skip this step if you have checked the medium before.
- 3 On the *Installation Mode* screen, choose *Update*. Clicking on *Next* will start the update procedure.

7.5.3 Updating via Subscription Management Tool (SMT)

As an alternative to downloading the updates for each single client system from the Novell update server, it is possible to use the Subscription Management Tool (SMT) for SUSE Linux Enterprise to mirror the updates to a local server.

This tool acts as Novell Customer Center proxy both for client registrations and as software update repository. The SMT documentation at <http://www.suse.com/doc/smt11/> gives an overview of its features as well as instructions on how to implement it.

7.5.4 Updating via SUSE Manager

SUSE Manager is a server solution for providing updates, patches, and security fixes for SUSE Linux Enterprise clients. It comes with a set of tools and a Web-based user interface for management tasks.

The SUSE Manager documentation at http://www.suse.com/doc/suse_manager/ gives an overview of its features as well as instructions on how to set up server and clients.

7.6 Updating SLE 11 SP2 to SLE 11 SP3

Online Migration is supported by the following tools:

- *YaST wagon* (graphical user interface)
- *zypper* (command line)

If updating your system via online migration, the update is carried out while the system is running. You only need to reboot once, after the update is completed. It is still possible to update with the following alternatives:

- Section 7.5.2, “Updating by Booting from an Installation Source” (page 152)

- Section 7.5.3, “Updating via Subscription Management Tool (SMT)” (page 155)
- Section 7.5.4, “Updating via SUSE Manager” (page 155)

7.6.1 Requirements

In order to do an online update, the following requirements must be met. Make sure to also read Section 7.4, “General Preparations for Updating” (page 145).

Product Registration

In order to be able to connect to the update channels, your product has to be registered. If this is not the case, either run the *Novell Customer Center Configuration* module in YaST or the `suse_register` command line tool to start the registration.

Run an Online Update

Make sure the currently installed version has the latest patches installed. Run an Online Update prior to the Online Migration. When using a graphical interface, start the YaST Online Update or the updater applet. On the command line, run the following commands (the last command needs to be run twice):

```
zypper ref -s
zypper update -t patch
zypper update -t patch
```

Reboot the system if needed.

See Chapter 1, *YaST Online Update* (↑*Administration Guide*) or Section “Updating Software with Zypper” (Chapter 6, *Managing Software with Command Line Tools*, ↑*Administration Guide*) for more information on the online update tools.

Third-Party Software

If your setup involves third-party software or add-on software, test this procedure on another machine to make sure that the dependencies are not broken by the update.

IMPORTANT: Always Run a Complete Online Migration

The online migration always has to be completed from beginning to end. If an online migration is interrupted in between, the system will be corrupted beyond recovery.

7.6.2 Online Migration with YaST *Wagon*

- 1 When all requirements are met (see Section 7.5.1.1, “Requirements” (page 147)), the update applet in the tray will display a message that a distribution upgrade is available. Click it to start YaST *Wagon*. Alternatively run `/usr/sbin/wagon` as `root` from the command line.
- 2 Confirm the *Welcome* dialog with *Next*.
- 3 If *Wagon* finds that the requirements are not met (required maintenance updates are available but not yet installed) it will do an automatic self-update which may require a reboot. Follow the on-screen instructions.
- 4 Choose the update method in the following dialog. Choose *Customer Center* to use the default setup (recommended).

Click *Custom URL* to manually choose the software channels used for the online migration. A list of channels will be displayed, providing the opportunity to manually enable, disable, add, or delete channels. Add the SP3 update source(s). This can either be the SP3 installation media or the SP3-POOL and SP3-Updates channels. Click *OK* to return to the *Update Method* dialog.

If you want to review changes to the channel setup caused by the update process, select *Check Automatic Repository Changes*.

Proceed with *Next*.

- 5 The system will be re-registered. During this process the SP3-POOL and SP3-Updates channels will be added to the system (see Section 7.2.3, “Channel Model” (page 140) for more information). Confirm the addition of the channels.
- 6 If you have selected *Check Automatic Repository Changes* in the *Update Method* dialog, the list of repositories will be displayed, providing the opportunity to manually enable, disable, add, or delete channels. Proceed with *OK* when finished.
- 7 The *Distribution Upgrade Settings* screen opens presenting a summary of the update configuration. The following sections are available:

Add-On Products

You may add SUSE Linux Enterprise Server add-on products or third-party products here.

Update Options

Lists the actions that will be performed during the update. You can choose whether to download all packages before installing them (default, recommended), or whether to download and install packages one by one.

Packages

Statistical overview of the update.

Backup

Set backup options.

Click *Next* and *Start The Update* to proceed.

IMPORTANT: Aborting the Online Migration

It is safe to abort the online migration on this screen *prior* to clicking *Start The Update* and on all previous screens. Click *Abort* to leave the update procedure and restore the system to the state it was in prior to starting YaST Wagon. Follow the instructions on screen and perform a re-registration before leaving Wagon to remove the SP2 channels from your system.

8 During the update procedure the following steps are executed:

8a Packages will be updated.

8b SuSEconfig will be run.

8c The system will be rebooted (press *OK*).

8d The newly updated system will be re-registered.

9 Your system has been successfully updated to Service Pack 3.

7.6.3 Online Migration with zypper

- 1** When all requirements are met (see Section 7.5.1.1, “Requirements” (page 147)), the “products” needed for the online migration

are added to `/etc/products.d`. Get a list of these products by running the following command:

```
zypper se -t product | grep -h -- "-migration" | cut -d'|' -f2
```

This command should at least return `SUSE_SLES-SP3-migration`. Depending on the scope of your installation, more products may be listed.

- 2 Install the migration products retrieved in the previous step with the command `zypper in -t product LIST_OF_PRODUCTS`, for example

```
zypper in -t product SUSE_SLES-SP3-migration
```

- 3 Register the products installed in the previous step in order to get the respective update channels:

```
suse_register -d 2 -L /root/.suse_register.log
```

- 4 Refresh the repositories and services:

```
zypper ref -s
```

- 5 Check the list of repositories you can retrieve with `zypper lr`.

If any of these repositories is not enabled (the SP3 ones are not enabled by default when following this workflow), enable them with `zypper modifyrepo --enable REPOSITORY ALIAS`, for example:

```
zypper modifyrepo --enable SLES11-SP3-Core SLES11-SP3-Updates
```

If your setup contains third-party repositories that may not be compatible with SP3, disable them with `zypper modifyrepo --disable REPOSITORY ALIAS`.

- 6 Now everything is in place to perform the distribution upgrade with `zypper dup --from REPO 1 --from REPO 2 . . .`. Make sure to list all needed repositories with `--from`, for example:

```
zypper dup --from SLES11-SP3-Pool --from SLES11-SP3-Updates
```

Confirm with `y` to start the upgrade.

- 7 Upon completion of the distribution upgrade from the previous step, run the following command:

```
zypper update -t patch
```

- 8 Now that the upgrade to SP3 has been completed, you need to re-register your product:

```
suse_register -d 2 -L /root/.suse_register.log
```

- 9 Lastly, reboot your system.

- 10 Your system has been successfully updated to Service Pack 3.

7.7 Backporting Source Code

SUSE uses backports extensively. The information in this section helps you understand why it can be deceptive to compare version numbers in order to judge software's capabilities and problems.

7.7.1 Why Backporting?

Upstream developers are primarily concerned with advancing the software they develop. In many cases they combine fixing bugs with introducing new features which have not yet received extensive testing and which may introduce new bugs.

For distribution developers, it is important to distinguish between:

- bugfixes with a limited potential for disrupting functionality; and
- changes that may disrupt existing functionality.

In most cases, distribution developers do not follow all upstream changes once a package has become part of a released distribution. Usually they stick instead with the upstream version that they initially released and create patches based on upstream changes to fix bugs. This practice is known as *backporting*.

Distribution developers generally will only introduce a newer version of software in two cases:

- when the changes between their packages and the upstream versions have become so large that backporting is no longer feasible, or
- for software that inherently ages badly, like anti-malware software.

7.7.2 Reasons for Backporting

SUSE uses backports extensively as we strike a good balance between a number of concerns for enterprise software. The most important of these are:

- Having stable interfaces (APIs) that software vendors can rely on when building products for use on SUSE's enterprise products.
- Ensuring that packages used in the release of SUSE's enterprise products are of the highest quality and have been thoroughly tested, both in themselves and as part of the whole enterprise product.
- Maintaining the various certifications of SUSE's enterprise products by other vendors, like certifications for Oracle or SAP products.
- Allowing SUSE's developers to focus on making the next version of the product as good as they can make it, rather than them having to spread their focus thinly across a wide range of releases.
- Keeping a clear view of what is in a particular enterprise release, so that our support can provide accurate and timely information about it.

7.7.3 Reasons against Backports

It is a general policy rule that no new upstream versions of a package are introduced into our enterprise products. This rule is not an absolute rule however. For a limited class of packages, in particular anti-virus software, security concerns weigh heavier than the conservative approach that is preferable from the perspective of quality assurance. For packages in that class, occasionally newer versions are introduced into a released version of an enterprise product line.

Sometimes also for other types of packages the choice is made to introduce a new version rather than a backport. This is done when producing a backport is not economically feasible or when there is a very relevant technical reason to introduce the newer version.

7.7.4 The Implications of Backports for Interpreting Version Numbers

Due to the practice of backporting, one cannot simply compare version numbers to determine whether a SUSE package contains a fix for a particular issue or has had a particular feature added to it. With backporting, the upstream part of a SUSE package's version number merely indicates what upstream version the SUSE package is based on. It may contain bug fixes and features that are not in the corresponding upstream release, but that have been backported into the SUSE package.

There are a number of locations where information regarding such bug fixes and features is stored:

- The package's changelog:

```
rpm -q --changelog name-of-installed-package
rpm -qp --changelog packagefile.rpm
```

The output briefly documents the change history of the package.

- The package changelog may contain entries like `bnc#1234` that refer to bugs in Novell's Bugzilla tracking system or links to other bugtracking systems. (Due to confidentiality policies, not all such information may be accessible to you).
- A package may contain a `/usr/share/doc/packagename/README.SUSE` or `README.SuSE` file which contains general, high-level information specific to the SUSE package.
- The RPM source package contains the patches that were applied during the building of the regular binary RPMs as separate files that can be interpreted if you are familiar with reading source code. See the Maximum RPM [<http://www.rpm.org/max-rpm/>] book for more information.
- For security bug fixes, the SUSE security announcements [<https://www.suse.com/support/security/#1>]. These often refer to bugs through standardized names like `CAN-2005-2495` which are maintained by the Common Vulnerabilities and Exposures [<http://cve.mitre.org>] project.

One particular area where this limited value of version numbers when backporting is involved can cause problems is with security scanning tools. Some security vulnerability scanning tools (or particular tests in such tools) operate solely on

version information. These tools/tests are thus prone to generating “false positives” (claims that a vulnerable piece of software has been found which in fact isn't vulnerable) when backports are involved. When evaluating reports from security scanning tools, one should always investigate whether an entry is based just on a version number or on an actual test of whether an actual vulnerability exists.

7.8 The Atomic Update

The Atomic Update is based on tools that manage two copies of the system and allow easy recovery after an update failure. The delivered tools require a special disk partition setup. Every copy of the system resides on a primary partition of its own. If an update fails, you can always switch back to the previous state of the system, which is available on the other partition.

7.8.1 Setup

WARNING: Strict Partitioning Requirements

The implementation has strict requirements on disk partitioning: the first root partition is `/dev/sda1` and must not occupy more than half of the entire disk size. Then the tools will create `/dev/sda2` for the system's second root partition. Further partitions if available are shared on both root partitions—take their size into account, and reduce the size of the first partition accordingly; this is a rough calculation:

The size of the complete disk minus size of `sda1` minus `sda2` is the free space of additional partitions.

- 1 Install the system with `/dev/sda1` as the single root partition and with less than half of the entire disk size.
- 2 Customize the installed system as needed. Make sure to have the `multi-update-tools` package installed.
- 3 Run `multi-update-setup --partition`, which creates the system's second root partition (`/dev/sda2`) of the similar size.
- 4 Partition the rest of the disk as needed and continue with customizations(*).

- 5 Run `multi-update-setup --clone` to copy the system to the other partition. With this command you also change the `/` (root) entry in `/etc/fstab` of the target system.
- 6 If needed, do further customizations(*).
- 7 Run `multi-update-setup --bootloader` to initialize the boot loader setup. The boot loader menu will then contain an entry to boot the other system.

WARNING: GRUB Bootloader Mandatory

Installation of the GRUB boot loader is mandatory. The tools are not compatible with other boot loaders.

- 8 If there are no customizations to be done as marked with (*), run `multi-update-setup --complete` which performs all the three steps.

7.8.2 Updating the Other System

Run `multi-update`. This command runs `zypper` in a `chroot` environment and updates the other system—it does not matter which one is active. Its boot menu will be offered as the default at boot time.

7.8.3 Troubleshooting

If the updated system has a damaged boot loader after the update, you must change the “Active” flag and set it for the root partition of the other system in order to boot it.

If the updated system does not boot at all, you need access to the boot loader menu to select the other system.

For more information about GRUB, see Chapter 10, *The Boot Loader GRUB* (↑*Administration Guide*).

7.8.4 Limitation

The root partition must be mounted by partition name, by ID, or in another way. Mounting by partition UUID or by label is not supported.

7.8.5 For More Information

For more information, see `/usr/share/doc/packages/multi-update-tools/README` coming with the `multi-update-tools` package.

7.9 Migration Hooks for YaST Wagon

Migration hooks allow you to run a custom external script at some point during the migration process. These scripts allow you to handle specific problems that cannot be handled via the usual RPM scripts, or allow you to perform any extra actions that might be needed during migration (not required during normal package update).

The migration hooks are executed with root privileges so it is possible to do any maintenance tasks in the scripts (starting/stopping services, data backup, data migration, etc...). The scripts must not be interactive; STDIN and STDOUT are redirected to pipes when running in YaST. The X session should not be used, as it might not be available in all cases (e.g. when running in text mode). Do not forget to set the executable permission for the hook scripts.

Migration hooks are supported in `yast2-wagon` package version 2.17.32.1 (provided as an update for SLES11-SP2) or 2.17.34 (included in SLES11-SP3) or higher versions.

7.9.1 Hook Script Location and Name Conventions

The scripts are searched in `/var/lib/YaST2/wagon/hooks/` directory. The expected script name is in the format `step_seq_prefix_name`, where:

step

is a predefined migration step name, describing the current migration step.

seq

is a sequence number in range 00...99, which makes it possible to set the order in which the scripts are executed. (It is important to keep the beginning zeros for correct sorting!)

prefix

should be unique to avoid conflicts (like a namespace). Use package name (if it is part of a package) or your vendor name, Internet domain name, etc., basically anything that can be considered sufficiently unique

name

can be any string (just to differentiate the scripts). Some descriptive name is recommended.

Example 7.2: *Hook Script with Full Path*

```
/var/lib/YaST2/wagon/hooks/before_package_migration_00_postgresql_backup
```

7.9.2 Hook Script Exit Value

The script should return exit value 0. If it fails (any non-zero exit value) an error message is displayed in Wagon and it is possible to restart the script, ignore the failure (and continue with other scripts) or completely cancel the hooks for the current step and stage.

7.9.3 Idempotent Scripts

The hook scripts *can be potentially run more times* (when going back and forth in the Wagon dialogs, Wagon might restart itself or some steps might be executed multiple times in the migration workflow), so the scripts have to cope with that fact (they can check at the beginning whether they need to do the action or the action has been already done or they can create a simple temporary stamp file or otherwise solve multiple runs properly).

7.9.4 List of Supported Hooks

Some hooks are optional (because they depend on the previous results or depend on user selected values). Note that some hooks are called multiple times (e.g. registration is called before migration and after migration). Here is the list of supported hooks (step names) in execution order:

`before_init`

started at the very beginning (note: it is called again after Wagon restarts)

`before_welcome , after_welcome`
started before/after displaying the welcome dialog

`before_registration_check , after_registration_check`
Wagon checks the registration status (if registration of some products has expired the migration might fail). If everything is OK, no dialog is displayed and Wagon automatically continues with the next step

`before_custom_url , after_custom_url`
repository manager is started (optional, in Patch CD mode only)

`before_self_update , after_self_update`
called before/after Wagon updates itself (to ensure the latest version is used for migration)

`before_installing_migration_products ,
after_installing_migration_products`
called before/after installing the migration products

`before_selecting_migration_source ,
after_selecting_migration_source`
Wagon asks the user to migrate via Novell Customer Center repositories or using a custom repository; the next step depends on the user selection

`before_registration , after_registration`
running SUSE register (to add migration repositories)

`before_repo_selection , after_repo_selection`
manual repository management

`before_set_migration_repo , after_set_migration_repo`
selecting migration repositories (full/minimal migration when using Novell Customer Center) or update repository selection (custom repository migration)

`before_package_migration`
before package update starts, after this step the real migration starts and it is not possible to go back to the previous state automatically (aborting in this phase results in an inconsistent (half upgraded) system, and manual rollback is needed)

`before_registration , after_registration`
running SUSE register (to register updated products)

`before_congratulate` , `after_congratulate`

before/after Wagon displays the congratulation dialog as a result of a successful migration

`before_exit`

called just before Wagon exits (always, regardless the migration result, also after abort and at restart)

7.9.5 Abort Hooks

These are special abort hooks which are called when the user aborts the migration. These hooks can be called in any step in the migration workflow therefore the execution order cannot be guaranteed. The scripts need to check the current state if they rely on the results of other hooks.

`before_abort`

user confirmed aborting the migration

`before_abort_rollback` , `after_abort_rollback`

user confirmed rollback after abort (reverting to the old products installed before starting migration). These hooks are called after `before_abort` and skipped when the user does not confirm rollback.

7.9.6 Restart Hooks

These hooks are called whenever Wagon restarts itself.

`before_restart`

Wagon is finishing and will be started again

`after_restart`

Wagon has restarted and runs the next step in the migration workflow

7.9.7 Usually Used Hooks

The list of hooks is fairly large, but many of them only make sense in special cases. In normal use cases these should be given preference:

- To do some action before the system is migrated (still running the previous version) use the `before_package_migration` hook.

At this point it is clear that the migration is ready and is about to start, whereas in all steps before it was possible to abort the migration.

- To do some action after the system has migrated (the system is running the new migrated version, but some things might not be active yet, e.g. updated kernel requires reboot, updated services might need restart etc.), use `before_congratulate` or `after_congratulate` hook.

This can be also used for cleaning up the temporary results of the `before_package_migration` hook. At this point the migration has successfully finished.

- To reverse the changes if the migration is aborted, use one of the abort hooks depending on the case. Keep in mind that the abort hooks can be called anytime, so the revert might not be needed (the hook that does the changes might not have been called yet). The abort hooks need to check the current state.

7.9.8 Obsolete Hooks

Older versions of Wagon supported only two hook scripts: `/usr/lib/YaST2/bin/wagon_hook_init` and `/usr/lib/YaST2/bin/wagon_hook_finish`. The problem was that only one script could be run as a hook and it was not possible to put hooks directly into RPM packages.

These old hook scripts are still supported in newer versions of Wagon for backward compatibility, but the new hooks `before_init` and `before_exit` should be used instead of the obsolete ones.

Setting Up Hardware Components with YaST

YaST allows you to configure hardware items at installation time as well as on an already-installed system. Configure audio hardware, printers or scanner support or learn which hardware components are connected to your computer by using the YaST Hardware Information module.

TIP: Graphics card, monitor, mouse and keyboard settings

Graphics card, monitor, mouse and keyboard can be configured with either KDE or GNOME tools.

8.1 Hardware Information

Use the YaST hardware information module if you want to know more about your hardware or if you need to find out details like vendor and model of a certain piece of hardware to be able to properly configure it.

- 1 Start YaST and click *Hardware > Hardware Information*. Hardware probing starts immediately and it will take some time until you see the hardware information tree in a separate window.
- 2 In the hardware information tree recursively click on the plus icons to expand the information about a specific device.
- 3 Click *Save to File...* to save the output to a file.
- 4 Click *Close* to leave the hardware information overview.

8.2 Setting Up Graphics Card and Monitor

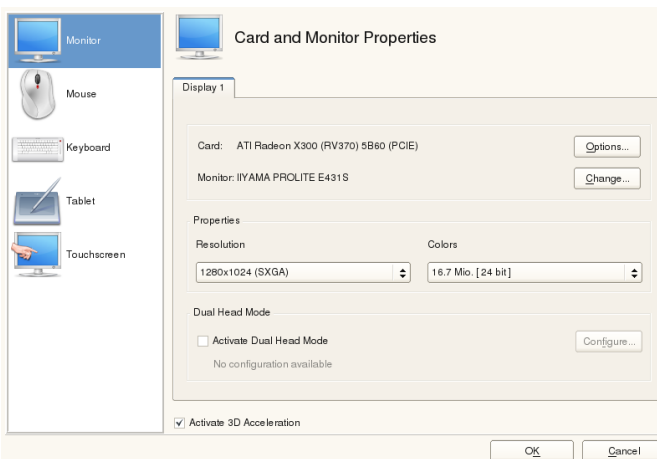
After the installation you can change the configuration of your graphics system (graphics card and monitor) according to your needs. Such a change may be necessary because of accessibility issues or hardware upgrades.

WARNING: Changing Monitor Frequencies

Although there are safety mechanisms, you should still be very careful when manually changing the allowed monitor frequencies. Incorrect values can damage your monitor. Always refer to the monitor's manual before changing frequencies.

Change the resolution, if fonts are too small or if circles appear misshapen. Proceed as follows:

- 1 In YaST, click *Hardware > Graphics Card and Monitor*. SaX2 checks the system resources and displays a window.
- 2 Make sure the monitor is properly detected. If not, use *Change* to select the appropriate model from the list.
- 3 Select an appropriate *Resolution* and *Colors*, if necessary.



- 4 Test the new configuration before it is applied to the system. Click *Ok* to decide what to do with your configuration (*Test*, *Save*, or *Cancel*.)

To activate a second monitor, proceed as follows:

- 1 In YaST, click *Hardware > Graphics Card and Monitor*. SaX2 checks the system resources and displays the *Card and Monitor Properties* dialog.
- 2 Make sure the monitor is properly detected. If not, use *Change* to select the appropriate model from the list.
- 3 Enable *Activate Dual Head Mode* and click *Configure* for further fine-tuning.
- 4 Make sure the second monitor is properly detected. If not, use *Change* to select the appropriate model from the list.
- 5 Decide whether you want to use the second monitor in *Cloned Multihead* or in *Xinerama Multihead* mode and click *Ok*.
- 6 Test the new configuration before it is applied to the system. Click *Ok* to decide what to do with your configuration (*Test*, *Save*, or *Cancel*.)

NOTE: Restarting the X Server

Any changes you make here take effect only after you restart the X server. If you want to restart the X server now, log out of the graphical system and log in again.

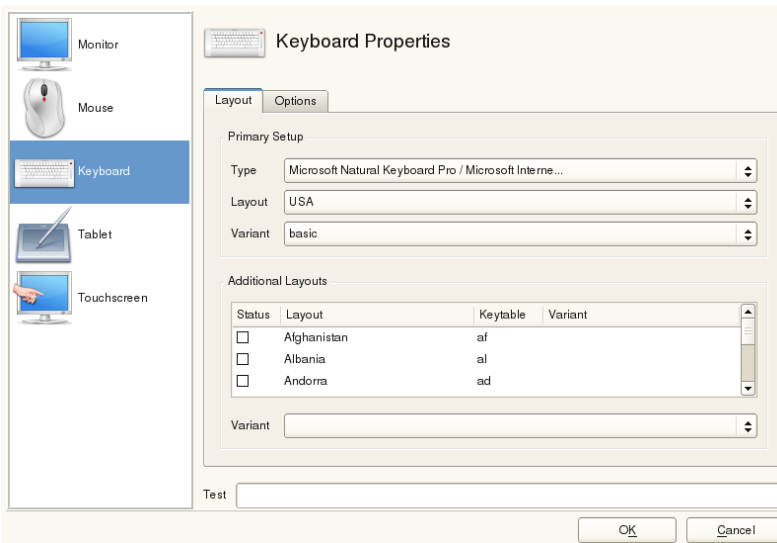
8.3 Setting Up Keyboard and Mouse

Reconfigure input devices such as the keyboard or the mouse, or add more than one of these devices using the YaST Keyboard and Mouse modules.

8.3.1 Keyboard Layout

In case you want to replace a standard 104-key keyboard with a multimedia keyboard or use a different language or country layout, proceed as follows:

- 1 In YaST, click *Hardware > Keyboard Layout*. The SaX2 configuration tool reads the system resources and displays the *Keyboard Properties* dialog.



- 2 Select your keyboard model from the *Type* list.
- 3 Select the country in the *Layout* list.
- 4 Depending on the country layout, you can choose a certain *Variant*. The selections are applied immediately for testing.
- 5 As an option you can enable *Additional Layouts*. Check one or more boxes in the list. This feature is handy if you want to switch between different languages or scripts in the running system without the need for reconfiguration.
- 6 Before saving the configuration, use the *Test* field at the bottom of the dialog to check if special characters like umlauts and accented characters can be entered and displayed correctly.
- 7 Click *OK* to leave the configuration dialog and in the following message click *Save* to apply your changes.

NOTE: Configuring Console Keyboard Layout

By clicking the *Save* button as described in Step 7 (page 174) the setup of the console keyboard layout takes place at the same time. If you want to change the console keyboard layout, either call `yast keyboard` (the text

mode interface) or check the `KEYTABLE` and `YAST_KEYBOARD` settings in `/etc/sysconfig/keyboard`.

8.3.2 Mouse Model

The mouse is usually detected automatically, but you can set up your mouse model manually if the automatic detection fails. Refer to the documentation of your mouse for a description of the model. If you want to modify your mouse configuration, proceed as follows:

- 1 In YaST, click *Hardware > Mouse Model*. The SaX2 configuration tool reads the system resources and displays the *Mouse Properties* dialog.
- 2 Click *Change* and select your mouse model from the list displayed.
- 3 Click *OK* to leave the configuration dialog and apply your changes with *Save*.

In the *Options* part of the dialog, set various options for operating your mouse.

Activate 3-Button Emulation

If your mouse has only two buttons, a third button is emulated whenever you click both buttons simultaneously.

Activate Mouse Wheel

Check this box to use a scroll wheel.

Invert X-Axis / Invert Y-Axis

Check these options if you want to change the direction in which the mouse pointer moves.

Activate Left-Hand Button Mapping

Check this box to make the button mapping suitable for left-hand usage.

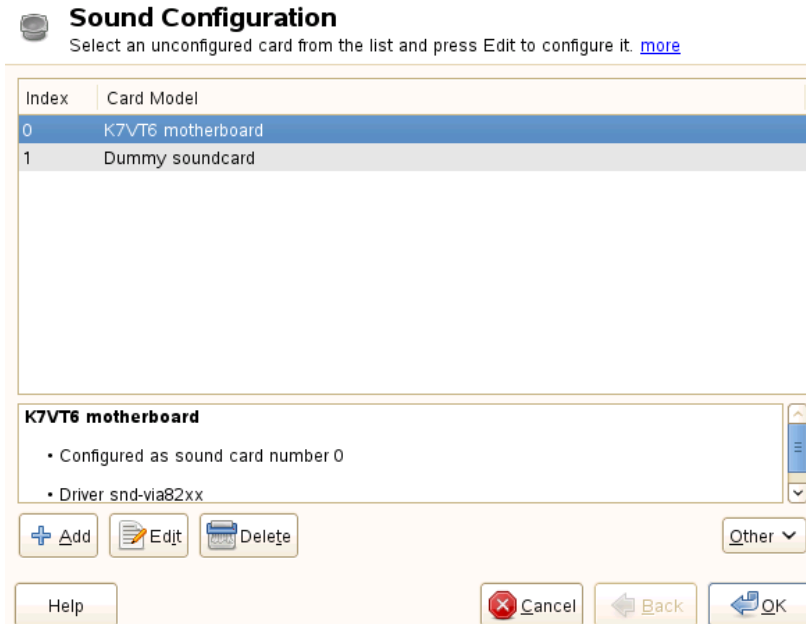
Emulate Wheel with Mouse Button

If your mouse does not have a scroll wheel but you would like to use a similar functionality, you can assign an additional button for this. Select the button to use. While pressing this button, any movement of the mouse is translated into scroll wheel commands. This feature is especially useful with trackballs.

8.4 Setting Up Sound Cards

YaST detects most sound cards automatically and configures them with the appropriate values. If you want to change the default settings, or need to set up a sound card that could not be configured automatically, use the YaST sound module. There, you can also set up additional sound cards or switch their order.

To start the sound module, start YaST and click *Hardware > Sound*. Alternatively, start the *Sound Configuration* dialog directly by running `yast2 sound` & as user `root` from a command line.



The dialog shows all sound cards that could be detected.

Procedure 8.1: *Configuring Sound Cards*

If you have added a new sound card or YaST could not automatically configure an existing sound card, follow the steps below. For configuring a new sound card, you need to know your sound card vendor and model. If in doubt, refer to your sound card documentation for the required information. For a reference list of sound cards supported by ALSA with their corresponding sound modules, see <http://www.alsa-project.org/main/index.php/Matrix:Main>.

During configuration, you can choose between the following setup options:

Quick Automatic Setup

You are not required to go through any of the further configuration steps—the sound card is configured automatically. You can set the volume or any options you want to change later.

Normal Setup

Allows you to adjust the output volume and play a test sound during the configuration.

Advanced setup with possibility to change options

For experts only. Allows you to customize all parameters of the sound card.

IMPORTANT: Advanced Configuration

Only use this option if you know exactly what you are doing. Otherwise leave the parameters untouched and use the normal or the automatic setup options.

- 1** Start the YaST sound module.
- 2** To configure a detected, but *Not Configured* sound card, select the respective entry from the list and click *Edit*.

To configure a new sound card, click *Add*. Select your sound card vendor and model and click *Next*.
- 3** Choose one of the setup options and click *Next*.
- 4** If you have chosen *Normal Setup*, you can now *Test* your sound configuration and make adjustments to the volume. You should start at about ten percent volume to avoid damage to your hearing or the speakers.
- 5** If all options are set according to your wishes, click *Next*.

The *Sound Configuration* dialog shows the newly configured or modified sound card.
- 6** To remove a sound card configuration that you no longer need, select the respective entry and click *Delete*.
- 7** Click *OK* to save the changes and leave the YaST sound module.

Procedure 8.2: *Modifying Sound Card Configurations*

- 1 To change the configuration of an individual sound card (for experts only!), select the sound card entry in the *Sound Configuration* dialog and click *Edit*.

This takes you to the *Sound Card Advanced Options* where you can fine-tune a number of parameters. For more information, click *Help*.

- 2 To adjust the volume of an already configured sound card or to test the sound card, select the sound card entry in the *Sound Configuration* dialog and click *Other*. Select the respective menu item.

NOTE: YaST Mixer

The YaST mixer settings provide only basic options. They are intended for troubleshooting (for example, if the test sound is not audible). Access the YaST mixer settings from *Other > Volume*. For everyday use and fine-tuning of sound options, use the mixer applet provided by your desktop or the `alsasound` command line tool.

- 3 For playback of MIDI files, select *Other > Start Sequencer*.
- 4 When a supported sound card is detected (like a Creative Soundblaster Live, Audigy or AWE sound card), you can also install SoundFonts for playback of MIDI files:
 - 4a Insert the original driver CD-ROM into your CD or DVD drive.
 - 4b Select *Other > Install SoundFonts* to copy SF2 SoundFonts™ to your hard disk. The SoundFonts are saved in the directory `/usr/share/sfbank/creative/`.
- 5 If you have configured more than one sound card in your system you can adjust the order of your sound cards. To set a sound card as primary device, select the sound card in the *Sound Configuration* and click *Other > Set as the Primary Card*. The sound device with index 0 is the default device and thus used by the system and the applications.
- 6 Per default, SUSE Linux Enterprise Server uses the PulseAudio sound system. It is an abstraction layer that helps to mix multiple audio streams, bypassing any restrictions the hardware may have. To enable or disable the PulseAudio

sound system, click *Other > PulseAudio Configuration*. If enabled, PulseAudio daemon is used to play sounds. Disable *PulseAudio Support* in case you want to use something else system-wide.

The volume and configuration of all sound cards are saved when you click *OK* and leave the YaST sound module. The mixer settings are saved to the file `/etc/asound.state`. The ALSA configuration data is appended to the end of the file `/etc/modprobe.d/sound` and written to `/etc/sysconfig/sound`.

8.5 Setting Up a Printer

YaST can be used to configure a local printer that is directly connected to your machine (normally with USB or parallel port) and to set up printing with network printers. It is also possible to share printers over the network. Further information about printing (general information, technical details, and troubleshooting) is available in Chapter 13, *Printer Operation* (*Administration Guide*).

In YaST, click *Hardware > Printer* to start the printer module. By default it opens in the *Printer Configurations* view, displaying a list of all printers that are available and configured. This is especially useful when having access to a lot of printers via the network. From here you can also *Print a Test Page* and configure local printers.

8.5.1 Configuring Local Printers

Usually a local USB printer is automatically detected. There are two possible reasons why a USB printer is not automatically detected:

- The USB printer is switched off.
- The communication between printer and computer is not possible. Check the cable and the plugs to make sure that the printer is properly connected. If this is the case, the problem may not be printer-related, but rather a USB-related problem.

Configuring a printer is basically a three-step process: specify the connection type, choose a driver, and name the printing queue for this setup.

For many printer models, several drivers are available. When configuring the printer, YaST defaults to the one marked `recommended` as a general rule. Normally it is

not necessary to change the driver—the recommended one should produce the best results. However, if you want a color printer to print only in black and white, it is most convenient to use a driver that does not support color printing, for example. If you experience performance problems with a PostScript printer when printing graphics, it may help to switch from a PostScript driver to a PCL driver (provided your printer understands PCL).

If no driver for your printer is listed, try to select a generic driver with an appropriate standard language from the list. Refer to your printer's documentation to find out which language (the set of commands controlling the printer) your printer understands. If this does not work, refer to Section 8.5.1.1, “Adding Drivers with YaST” (page 181) for another possible solution.

A printer is never used directly, but always through a print queue. This ensures that simultaneous jobs can be queued and processed one after the other. Each print queue is assigned to a specific driver, and a printer can have multiple queues. This makes it possible to set up a second queue on a color printer that prints black and white only, for example. Refer to Section “The Workflow of the Printing System” (Chapter 13, *Printer Operation*, ↑*Administration Guide*) for more information about print queues.

Procedure 8.3: *Adding a New Local Printer*

- 1** Start the YaST printer module with *Hardware > Printer*.
- 2** In the *Printer Configurations* screen click *Add*.
- 3** If your printer is already listed under *Specify the Connection*, proceed with the next step. Otherwise, try to *Detect More* or start the *Connection Wizard*.
- 4** In the input box under *Find and Assign a Driver* enter the vendor name and the model name and click *Search for*.
- 5** Choose the driver marked as recommended that best matches your printer. If no suitable driver is displayed:
 - 5a** Check your search term
 - 5b** Broaden your search by clicking *Find More*
 - 5c** Add a driver as described in Section 8.5.1.1, “Adding Drivers with YaST” (page 181)

- 6 Specify the `Default paper size`.
- 7 In the *Set Arbitrary Name* field, enter a unique name for the print queue.
- 8 The printer is now configured with the default settings and ready to use. Click *OK* to return to the *Printer Configurations* view. The newly configured printer is now visible in the list of printers.

8.5.1.1 Adding Drivers with YaST

If no suitable driver is available in the *Find and Assign a Driver* dialog when adding a new printer, no PPD (PostScript Printer Description) file for your model is available. For more information about PPD files, refer to Section “Installing the Software” (Chapter 13, *Printer Operation*, ↑*Administration Guide*).

Get PPD files directly from your printer vendor or from the driver CD of a PostScript printer. For details, see Section “No Suitable PPD File Available for a PostScript Printer” (Chapter 13, *Printer Operation*, ↑*Administration Guide*). Alternatively, find PPD files at <http://www.linuxfoundation.org/collaborate/workgroups/openprinting/database/databaseintro>, the “OpenPrinting.org printer database”. When downloading PPD files from OpenPrinting, keep in mind that it always shows the latest Linux support status, which is not necessarily met by SUSE Linux Enterprise Server.

Procedure 8.4: *Adding a PPD file*

- 1 Start the YaST printer module with *Hardware > Printer*.
- 2 In the *Printer Configurations* screen, click *Add*.
- 3 In the *Find and Assign a Driver* section, click *Driver Packages*.
- 4 Enter the full path to the PPD file into the input box under *Make a Printer Description File Available*. Alternatively, choose the file from a dialog box by clicking *Browse*.
- 5 Click *OK* to return to the *Add New Printer Configuration* screen.
- 6 In order to directly use this PPD file, proceed as described in Procedure 8.3, “Adding a New Local Printer” (page 180). Otherwise, click *Cancel*.

8.5.1.2 Editing a Local Printer Configuration

By editing an existing configuration for a local printer you cannot only change basic settings as connection type and driver, but also adjust the default settings for paper size, resolution, media source, etc. You can change the identifier of the printer by altering the printer descriptions.

Procedure 8.5: *Editing a Local Printer*

- 1 Start the YaST printer module with *Hardware > Printer*.
- 2 In the *Printer Configurations* screen, choose a local printer from the list and click *Edit*.
- 3 Change the connection type or the driver as described in Procedure 8.3, “Adding a New Local Printer” (page 180). This should only be necessary in case you have problems with the current configuration.
- 4 Make this printer the default by checking *Default Printer*.
- 5 Adjust the default settings by clicking *All Options for the Current Driver*. To change a setting, expand the list of options by clicking the relative + sign. Change the default by clicking an option. Apply your changes with *OK*.

8.5.2 Configuring Printing via the Network with YaST

Network printers are not detected automatically. They must be configured manually using the YaST printer module. Depending on your network setup, you can print to a print server (CUPS, LPD, SMB, or IPX) or directly to a network printer (preferably via TCP). Access the configuration view for network printing by choosing *Printing via Network* from the left pane in the YaST printer module.

8.5.2.1 Using CUPS

In a Linux environment CUPS is usually used to print via the network. The simplest setup is to only print via a single CUPS server which can directly be accessed by

all clients. Printing via more than one CUPS server requires a running local CUPS daemon that communicates with the remote CUPS servers.

Procedure 8.6: *Printing via a Single CUPS Server*

- 1 Start the YaST printer module with *Hardware > Printer*.
- 2 From the left pane, launch the *Print via Network* screen.
- 3 Check *Do All Your Printing Directly via One Single CUPS Server* and specify the name or IP address of the server.
- 4 Click *Test Server* to make sure you have chosen the correct name or IP address.
- 5 Click OK to return to the *Printer Configurations* screen. All printers available via the CUPS server are now listed.

Procedure 8.7: *Printing via Multiple CUPS Servers*

- 1 Start the YaST printer module with *Hardware > Printer*.
- 2 From the left pane, launch the *Print via Network* screen.
- 3 Check *Accept Printer Information from the Following Servers*
- 4 Under *General Settings* specify which servers to use. You may accept connections from all networks available, from the local network, or from specific hosts. If you choose the latter option, you need to specify the hostnames or IP addresses.
- 5 Confirm by clicking *OK* and then *Yes* when asked to start a local CUPS server. After the server has started YaST will return to the *Printer Configurations* screen. Click *Refresh list* to see the printers detected by now. Click this button again, in case more printer are to be available.

8.5.2.2 Using Print Servers other than CUPS

If your network offers print services via print servers other than CUPS, start the YaST printer module with *Hardware > Printer* and launch the *Print via Network* screen from the left pane. Start the *Connection Wizard* and choose the appropriate

Connection Type. Ask your network administrator for details on configuring a network printer in your environment.

8.5.3 Sharing Printers Over the Network

Printers managed by a local CUPS daemon can be shared over the network and so turn your machine into a CUPS server. Usually you share a printer by enabling CUPS' so-called “browsing mode”. If browsing is enabled, the local print queues are made available on the network for listening to remote CUPS daemons. It is also possible to set up a dedicated CUPS server that manages all printing queues and can directly be accessed by remote clients. In this case it is not necessary to enable browsing.

Procedure 8.8: *Sharing Printers*

- 1** Start the YaST printer module with *Hardware > Printer*.
- 2** Launch the *Share Printers* screen from the left pane.
- 3** Select *Allow Remote Access*. For more detailed configuration, additional options are available:
 - Check *For computers within the local network* and enable browsing mode by also checking *Publish printers by default within the local network*.
 - *Add* the network interface to be used by the CUPS server. If you want to share your printers via specified network interfaces, add those in the input box below.
 - To restrict access to your CUPS server to certain networks or IP addresses, specify these via the two input boxes.
- 4** Click *OK* to restart the CUPS server and to return to the *Printer Configurations* screen.
- 5** Regarding CUPS and firewall settings, see http://en.opensuse.org/SDB:CUPS_and_SANE_Firewall_settings.

8.6 Setting Up a Scanner

You can configure a USB or SCSI scanner with YaST. The `sane-backends` package contains hardware drivers and other essentials needed to use a scanner. Scanners connected to a parallel port cannot be configured with YaST. If you own a HP All-In-One device, see Section 8.6.1, “Configuring an HP All-In-One Device” (page 185), instructions on how to configure a network scanner are available at Section 8.6.3, “Scanning over the Network” (page 186).

Procedure 8.9: *Configuring a USB or SCSI Scanner*

- 1 Connect your USB or SCSI scanner to your computer and turn it on.
- 2 Start YaST and select *Hardware > Scanner*. YaST builds the scanner database and tries to detect your scanner model automatically.

If a USB or SCSI scanner is not properly detected, try *Other > Restart Detection*.

- 3 To activate the scanner select it from the list of detected scanners and click *Edit*.
- 4 Choose your model form the list and click *Next* and *Finish*.
- 5 Use *Other > Test* to make sure you have chosen the correct driver.
- 6 Leave the configuration screen with *OK*.

8.6.1 Configuring an HP All-In-One Device

An HP All-In-One device can be configured with YaST even if it is connected to the parallel port or is made available via the network. If you own a USB HP All-In-One device, start configuring as described in Procedure 8.9, “Configuring a USB or SCSI Scanner” (page 185). If it is detected properly and the *Test* succeeds, it is ready to use.

If your USB device is not properly detected, or your HP All-In-One device is connected to the parallel port or the network, run the HP Device Manager:

- 1 Start YaST and select *Hardware > Scanner*. YaST loads the scanner database.
- 2 Start the HP Device Manager with *Other > Run hp-setup* and follow the on-screen instructions. After having finished the HP Device Manager, the YaST scanner module automatically restarts the auto detection.

- 3 Test it by choosing *Other > Test*.
- 4 Leave the configuration screen with *OK*.

8.6.2 Sharing a Scanner over the Network

SUSE Linux Enterprise Server allows the sharing of a scanner over the network. To do so, configure your scanner as follows:

- 1 Configure the scanner as described in Section 8.6, “Setting Up a Scanner” (page 184).
- 2 Choose *Other > Scanning via Network*.
- 3 Enter the hostnames of the clients (separated by a comma) that should be allowed to use the scanner under *Server Settings > Permitted Clients for saned* and leave the configuration dialog with *OK*.

8.6.3 Scanning over the Network

To use a scanner that is shared over the network, proceed as follows:

- 1 Start YaST and select *Hardware > Scanner*.
- 2 Open the network scanner configuration menu by *Other > Scanning via Network*.
- 3 Enter the hostname of the machine the scanner is connected to under *Client Settings > Servers Used for the net Metadriver*.
- 4 Leave with *OK*. The network scanner is now listed in the Scanner Configuration window and is ready to use.

Installing or Removing Software

Use YaST's software management tool to search for software components you want to add or remove. YaST resolves all dependencies for you. To install packages not shipped with the installation media, add additional software repositories to your setup and let YaST manage them. Keep your system up-to-date by managing software updates with the update applet.

Change the software collection of your system with the YaST Software Manager. This YaST module is available in two toolkit flavors: Qt and ncurses; the Qt flavor is described here.

NOTE: Confirmation and Review of Changes

When installing, updating or removing packages, any changes in the Software Manager are not applied immediately but only after confirming them with *Accept* or *Apply* respectively. YaST maintains a list with all actions, allowing you to review and modify your changes before applying them to the system.

9.1 Definition of Terms

Repository

A local or remote directory containing packages, plus additional information about these packages (package meta-data).

(Repository) Alias

A short name for a repository used by various Zypper commands. The alias can be chosen by the user when adding a repository and must be unique.

Product

Represents a whole product, for example SUSE® Linux Enterprise Server.

Pattern

A pattern is an installable group of packages dedicated to a certain purpose. For example, the `Laptop` pattern contains all packages that are needed in a mobile computing environment. Patterns define package dependencies (such as required or recommended packages) and come with a preselection of packages marked for installation. This ensures that the most important packages needed for a certain purpose are available on your system after installation of the pattern. However, not necessarily all packages in a pattern are preselected for installation and you can manually select or deselect packages within a pattern according to your needs and wishes.

Package

A package is a compressed file in `rpm` format that contains the files for a particular program.

Patch

A patch consists of one or more packages and may be applied by means of `deltarpm`s. It may also introduce dependencies to packages that are not installed yet.

Resolvable

An generic term for product, pattern, package or patch. The most commonly used type of resolvable is a package or a patch.

deltarpm

A `deltarpm` consists only of the binary diff between two defined versions of a package, and therefore has the smallest download size. Before being installed, the full RPM package is rebuilt on the local machine.


Package Dependencies

Certain packages are dependent on other packages, such as shared libraries. In other terms, a package may *require* other packages—if the required packages are not available, the package cannot be installed. In addition to dependencies (package requirements) that must be fulfilled, some packages *recommend* other packages. These recommended packages are only installed if they are actually

available, otherwise they are just ignored and the package recommending them is installed nevertheless.

9.2 Using the KDE Interface (Qt)

The YaST Qt interface is started by default when using the desktops KDE, icewm, and others. It is also used when invoking YaST from a remote terminal. Start the software manager from the YaST Control Center by choosing *Software > Software Management*.



9.2.1 Views for Searching Packages or Patterns

The YaST software manager can install packages or patterns from all currently enabled repositories. It offers different views and filters to make it easier to find the

software you are searching for. The *Search* view is the default view of the window. To change view, click *View* and select one of the following entries from the drop-down list. The selected view opens in a new tab.

Patterns

Lists all patterns available for installation on your system.

Package Groups

Lists all packages sorted by groups such as *Graphics*, *Programming*, or *Security*.

RPM Groups

Lists all packages sorted by functionality with groups and subgroups. For example *Networking > Email > Clients*.

Languages

Filter to list all packages needed to add a new system language.

Repositories

Filter to list packages by repository. In order to select more than one repository, hold the Ctrl key while clicking on repository names. The “pseudo repository” *@System* lists all packages currently installed.

Search

Lets you search for a package according to certain criteria. Enter a search term and press Enter. Refine your search by specifying where to *Search In* and by changing the *Search Mode*. For example, if you do not know the package name but only the name of the application that you are searching for, try including the package *Description* in the search process.

Installation Summary

If you have already selected packages for installation, update or removal, this view shows the changes that will be applied to your system as soon as you click *Accept*. To filter for packages with a certain status in this view, activate or deactivate the respective check boxes. Hit Shift + F1 for details on the status flags.

TIP: Finding Packages not Belonging to an Active Repository

To list all packages that do not belong to an active repository, choose *View > Repositories > @System* and then choose *Secondary Filter > Unmaintained Packages*. This is useful, for example, if you have deleted

a repository and would like to make sure no packages from that repository remain installed.

9.2.2 Installing and Removing Packages or Patterns

Certain packages are dependent on other packages, such as shared libraries. On the other hand, some packages cannot coexist with others on the system. If possible, YaST automatically resolves these dependencies or conflicts. If your choice results in a dependency conflict that cannot be automatically solved, you need to solve it manually as described in Section 9.2.4, “Checking Software Dependencies” (page 194).

NOTE: Removal of Packages

When removing any packages, by default YaST only removes the selected packages. If you want YaST to also remove any other packages that become unneeded after removal of the specified package, select *Options > Cleanup when deleting packages*.

- 1 Search for packages as described in Section 9.2.1, “Views for Searching Packages or Patterns” (page 189).
- 2 The packages found are listed in the right pane. To install a package or remove it, right-click it and choose *Install* or *Delete*. If the relevant option is not available, check the package status indicated by the symbol in front of the package name—hit Shift + F1 for help.

TIP: Applying an Action to All Packages Listed

To apply an action to all packages listed in the right pane, choose an action from *Package > All in This List*.

- 3 To install a pattern, right-click the pattern name and choose *Install*.
- 4 It is not possible to remove a pattern per se. Instead, select the packages of a pattern you want to remove and mark them for removal.

- 5 In order to select more packages, repeat the steps mentioned above.
- 6 Before applying your changes, you can review or modify them by clicking *View > Installation Summary*. By default, all packages that will change status, are listed.
- 7 In order to revert the status for a package, right-click the package and select one of the following entries: *Keep* if the package was scheduled to be deleted or updated, or *Do Not Install* if it was scheduled for installation. To abandon all changes and close the Software Manager, click *Cancel* and *Abandon*
- 8 When you are finished, click *Accept* to apply your changes.
- 9 In case YaST found dependencies on other packages, a list of packages that have additionally been chosen for installation, update or removal is presented. Click *Continue* to accept them.

After all selected packages are installed, updated or removed, the YaST Software Manager automatically terminates.

NOTE: Installing Source Packages

Installing source packages with YaST Software Manager is not possible at the moment. Use the command line tool `zypper` for this purpose. For more information, see Section “Installing or Downloading Source Packages” (Chapter 6, *Managing Software with Command Line Tools*, ↑*Administration Guide*).

9.2.3 Updating Packages

Instead of updating individual packages, you can also update all installed packages or all packages from a certain repository. When mass updating packages, the following aspects are generally considered:

- priorities of the repositories that provide the package,
- architecture of the package (for example, x86_64, i686, i586),
- version number of the package,

- package vendor.

Which of the aspects has the highest importance for choosing the update candidates depends on the respective update option you choose.

- 1 To update all installed packages to the latest version, choose *Package > All Packages > Update if Newer Version Available* from the main menu.

All repositories are checked for possible update candidates, using the following policy: YaST first tries to restrict the search to packages with the same architecture and vendor like the installed one. If the search is positive, the “best” update candidate from those is selected according to the process below. However, if no comparable package of the same vendor can be found, the search is expanded to all packages with the same architecture. If still no comparable package can be found, all packages are considered and the “best” update candidate is selected according to the following criteria:

1. Repository priority: Prefer the package from the repository with the highest priority.
2. If more than one package results from this selection, choose the one with the “best” architecture (best choice: matching the architecture of the installed one; otherwise: x86_64 > i686 > i586).

If the resulting package has a higher version number than the installed one, the installed package will be updated and replaced with the selected update candidate.

This option tries to avoid changes in architecture and vendor for the installed packages, but under certain circumstances, they are tolerated.

NOTE: Update Unconditionally

If you choose *Package > All Packages > Update Unconditionally* instead, basically the same criteria apply but candidate package found is installed unconditionally. Thus, choosing this option might actually lead to downgrading some packages.

- 2 To make sure that the packages for a mass update derive from a certain repository:
 - 2a Choose the repository from which to update as described in Section 9.2.1, “Views for Searching Packages or Patterns” (page 189).

- 2b** On the right hand side of the window, click *Switch system packages to the versions in this repository*. This explicitly allows YaST to change the package vendor when replacing the packages.

As soon as you proceed with *Accept*, all installed packages will be replaced by packages deriving from this repository, if available. This may lead to changes in vendor and architecture and even to downgrading some packages.

- 2c** To refrain from this, click *Cancel switching system packages to the versions in this repository*. Note that you can only cancel this until you press the *Accept* button.

- 3** Before applying your changes, you can review or modify them by clicking *View > Installation Summary*. By default, all packages that will change status, are listed.
- 4** If all options are set according to your wishes, confirm your changes with *Accept* to start the mass update.

9.2.4 Checking Software Dependencies

Most packages are dependent on other packages. If a package, for example, uses a shared library, it is dependent on the package providing this library. On the other hand some packages cannot coexist with each other, causing a conflict (for example, you can only install one mail transfer agent: sendmail or postfix). When installing or removing software, the Software Manager makes sure no dependencies or conflicts remain unsolved to ensure system integrity.

In case there exists only one solution to resolve a dependency or a conflict, it is resolved automatically. Multiple solutions always cause a conflict which needs to be resolved manually. If solving a conflict involves a vendor or architecture change, it also needs to be solved manually. When clicking *Accept* to apply any changes in the Software Manager, you get an overview of all actions triggered by the automatic resolver which you need to confirm.

By default, dependencies are automatically checked. A check is performed every time you change a package status (for example, by marking a package for installation or removal). This is generally useful, but can become exhausting when manually resolving a dependency conflict. To disable this function, uncheck *Dependencies >*


Autocheck. Manually perform a dependency check with *Dependencies > Check Now*. A consistency check is always performed when you confirm your selection with *Accept*.

To review a package's dependencies, right-click it and choose *Show Solver Information*. A map showing the dependencies opens. Packages that are already installed are displayed in a green frame.

NOTE: Manually Solving Package Conflicts

Unless you are very experienced, follow the suggestions YaST makes when handling package conflicts, otherwise you may not be able to resolve them. Keep in mind that every change you make, potentially triggers other conflicts, so you can easily end up with a steadily increasing number of conflicts. In case this happens, *Cancel* the Software Manager, *Abandon* all your changes and start again.

Figure 9.1: *Conflict Management of the Software Manager*



9.3 Managing Software Repositories and Services

If you want to install third-party software, add additional software repositories to your system. By default, the product repositories such as SUSE Linux Enterprise Server-DVD 11 SP3 and a matching update repository are automatically configured once you have registered your system. For more information about registration, see Section 6.16.1.4, “Novell Customer Center Configuration” (page 127). Depending on the initially selected product, a separate language add-on repository with translations, dictionaries, etc. might also be configured.

To manage repositories, start YaST and select *Software > Software Repositories*. The *Configured Software Repositories* dialog opens. Here, you can also manage subscriptions to so-called *Services* by changing the *View* at the right corner of the dialog to *All Services*. A Service in this context is a *Repository Index Service* (RIS) that can offer one or more software repositories. Such a Service can be changed dynamically by its administrator or vendor.

Each repository provides files describing content of the repository (package names, versions, etc.). These repository description files are downloaded to a local cache that is used by YaST. To ensure their integrity, software repositories can be signed with the GPG Key of the repository maintainer. Whenever you add a new repository, YaST offers the ability to import its key.

WARNING: Trusting External Software Sources

Before adding external software repositories to your list of repositories, make sure this repository can be trusted. SUSE Linux Enterprise Server is not responsible for any potential problems arising from software installed from third-party software repositories.

9.3.1 Adding Software Repositories

You can either add repositories from a local hard disk, from a removable medium (like a CD, DVD or a USB mass storage) or from a network.

To add repositories from the *Configured Software Repositories* dialog in YaST proceed as follows:

1 Click *Add*.

2 From the list of available *Media Types* specify the type matching your repository:

For network sources, it is usually sufficient to use the default option, *Specify URL*.

To add a repository from a removable medium or a local hard disk, choose the relevant option and insert the medium or connect the USB device to the machine, respectively.

3 You can choose to *Download Repository Description Files* now. If the option is unchecked, YaST will automatically download the files later, if needed. Click *Next* to proceed.

4 When adding a repository from the network, enter the data you are prompted for. Continue with *Next*.

5 Depending on the repository you have added, you might be asked if you want to import the GPG key with which it is signed or asked to agree to a license.

After confirming these messages, YaST will download and parse the metadata and add the repository to the list of *Configured Repositories*.

6 If needed, adjust the repository *Properties* as described in Section 9.3.2, “Managing Repository Properties” (page 197) or confirm your changes with *OK* to close the configuration dialog.

Now you can install software from this repository as described in Section 9.2, “Using the KDE Interface (Qt)” (page 189).

9.3.2 Managing Repository Properties

The *Configured Software Repositories* overview of the *Software Repositories* lets you change the following repository properties:

Status

The repository status can either be *Enabled* or *Disabled*. You can only install packages from repositories that are enabled. To turn a repository off temporarily click *Disable*. You can also double-click a repository name to toggle its status. If you want to remove a repository completely, click *Delete*.

Refresh

When refreshing a repository, its content description (package names, versions, etc.) is downloaded to a local cache that is used by YaST. It is sufficient to do this once for static repositories such as CDs or DVDs, whereas repositories whose content changes often should be refreshed frequently. The easiest way to keep a repository's cache up-to-date is to choose *Automatically Refresh*. To do a manual refresh click *Refresh* and select one of the options.

Keep Downloaded Packages

Packages from remote repositories are downloaded before being installed. By default, they are deleted upon a successful installation. Activating *Keep Downloaded Packages* prevents the deletion of downloaded packages. The download location is configured in `/etc/zypp/zypp.conf`, by default it is `/var/cache/zypp/packages`.

Priority

The *Priority* of a repository is a value between 1 and 200, with 1 being the highest priority and 200 the lowest priority. Any new repositories that are added with YaST get a priority of 99 by default. If you do not care about a priority value for a certain repository, you can also set the value to 0 to apply the default priority to that repository (99). If a package is available in more than one repository, then the repository with the highest priority takes precedence. This is useful if you want to avoid downloading packages unnecessarily from the Internet by giving a local repository (for example, a DVD) a higher priority.

IMPORTANT: Priority vs. Version

The repository with the highest priority takes precedence in any case. Therefore, make sure that the update repository always has the highest priority (20 by default), otherwise you might install an outdated version that will not be updated until the next online update.

Name and URL

To change a repository name or its URL, select it from the list with a single-click and then click *Edit*.

9.3.3 Managing Repository Keys

To ensure their integrity, software repositories can be signed with the GPG Key of the repository maintainer. Whenever you add a new repository, YaST offers to

import its key. Verify it as you would do with any other GPG key and make sure it does not change. If you detect a key change, something might be wrong with the repository. Disable the repository as an installation source until you know the cause of the key change.

To manage all imported keys, click *GPG Keys...* in the *Configured Software Repositories* dialog. Select an entry with the mouse to show the key properties at the bottom of the window. *Add*, *Edit* or *Delete* keys with a click on the respective buttons.

9.4 Keeping the System Up-to-date

Novell offers a continuous stream of software security patches for your product. The update applet informs you about the availability of patches and lets you easily install them with just a few clicks.

9.4.1 Using the KDE Software Updater

The Software Updater icon resides in the system tray of your panel depicting a gearwheel with a green arrow. To start Software Updater manually, choose *System Settings > SoftwareManagement > Software Updates* from the main menu. Alternatively, **Alt + F2** and enter `kpk_update`.

NOTE: Icon Visibility


The Software Updater icon is only visible in the system tray, if patches are available. Hover over the icon to see the number of patches available.

9.4.1.1 Installing Patches

- 1 Whenever software updates are available, the applet icon appears in the panel. Left-click the Software Updater icon to launch the *Review and Update* software window.
- 2 Select a patch for installation by ticking its check box. Get detailed information on a patch by clicking on its title. To select all available patches for installation, tick the check box in the table header.
- 3 Click *Apply* to start the patch installation.

- 4 In case you have started the patch installation for the first time, you will be asked to enter the `root` password twice in order to proceed. If you also check *Remember authorization* you will never be asked again to provide the password.
- 5 The *Additional Changes* window showing an installation summary opens. Click *Continue* to finish the installation.

Figure 9.2: KDE Software Updater



The YaST Online Update offers advanced features to customize the patch installation. Please refer to Chapter 1, *YaST Online Update* (*Administration Guide*) for more information.

9.4.1.2 Configuring the KDE Software Updater

By default Software Updater checks for updates every 24 hours, notifies you when patches are available and does not automatically install patches. These settings can be changed with the *Software Management settings*. To open the *Software Management settings* choose *System Settings > Software Management > Settings* from the main menu. Alternatively, press `Alt + F2` and enter `kpk_settings`. The settings for Software Updater are available in the *Update Settings* section.

IMPORTANT: Patch Origin

The *Software Management settings* also allows you to configure the repositories (*Origin of Packages*) to be used. This setting not only applies

to Software Updater but also to the KDE Software Management module (*Get and Remove Software*).

Make sure the repository *Updates for SUSE Linux Enterprise Server 11 SP3* is always selected—otherwise you will not receive patches.

9.4.2 Using the GNOME Update Applet

The update applet resides in the notification area of the panel. Its icon changes depending on the availability and relevance of patches and the status of the update. To invoke the applet manually, choose *Computer > More Applications > System > Software Update*.

NOTE: Icon visibility

The applet icon is only visible if the following conditions are met:

- patches are available
- the GUI was not started as user `root`
- the GUI was not started in a VNC session

To start the update viewer even if no applet icon is visible, press `Alt + F2` and enter `gpk-update-viewer`.

Open box with a globe

The update applet is busy (for example checking for updates or installing software).

Red Star with Exclamation Mark

Security patches are available.

Orange Star with an Up Arrow

Important patches are available.

Yellow Star with a Down Arrow

Trivial patches are available.

Yellow Triangle with Exclamation Mark


An error has occurred.

9.4.2.1 Installing Patches

Procedure 9.1: *Installing Patches*

- 1 Whenever new patches are available, a notification message will appear and the Update Applet icon will be visible in the notification area. Either click *Install updates* in the notification message or click the icon to open the *Software Update* window.
- 2 All security and important patches are preselected. It is strongly recommended to install these patches. Trivial patches can be manually selected by ticking the respective check boxes. Get detailed information on a patch by clicking on its title.
- 3 Click *Install Updates* to start the patch installation.
- 4 The *Additional Confirmation Required* window showing an installation summary opens. Click *Continue* to proceed.
- 5 Enter the `root` password in the authentication screen and proceed with *Authenticate*.

Figure 9.3: *GNOME Update Applet*



The YaST Online Update offers advanced features to customize the patch installation. Please refer to Chapter 1, *YaST Online Update* (↑*Administration Guide*) for more information.

9.4.2.2 Configuring the Software Update Applet

To configure the update applet, right-click the update icon in the panel and choose *Preferences*. The configuration dialog lets you modify the following settings:

Check for Updates

Choose how often a check for updates is performed: *Hourly*, *Daily*, *Weekly*, or *Never*.

Automatically Install

Configure whether patches are installed automatically or not (default). Automatic installation can be chosen for either security patches only or for all patches.

Check for Major Upgrades

Choose how often a check for major upgrades is performed: *Daily*, *Weekly*, or *Never*.

Check for updates when using mobile broadband

This configuration option is only available on mobile computers. Turned off by default.

More options are configurable using `gconf-editor`: *apps > gnome-packagekit*.

Installing Add-On Products

Add-on products are system extensions. You can install a third party add-on product or a special system extension of SUSE® Linux Enterprise Server (for example, a CD with support for additional languages or a CD with binary drivers). To install a new add-on, start YaST and select *Software > Add-On Products* . You can select various types of product media, like CD, FTP, USB mass storage devices (such as USB flash drives or disks) or a local directory. You can work also directly with ISO files. To add an add-on as ISO file media, select *Local ISO Image* then enter the *Path to ISO Image*. The *Repository Name* is arbitrary.

10.1 Add-Ons

To install a new add-on, proceed as follows:

- 1 In YaST select *Software > Add-On Products* to see an overview of already installed add-on products.
- 2 To install a new add-on product, click *Add*.
- 3 From the list of available *Media Types* specify the type matching your repository.
- 4 To add a repository from a removable medium, choose the relevant option and insert the medium or connect the USB device to the machine, respectively.
- 5 You can choose to *Download Repository Description Files* now. If the option is unchecked, YaST will automatically download the files later, if needed. Click *Next* to proceed.

6 When adding a repository from the network, enter the data you are prompted for. Continue with *Next*.

7 Depending on the repository you have added, you may be asked if you want to import the GPG key with which it is signed or asked to agree to a license.

After confirming these messages, YaST will download and parse the metadata and add the repository to the list of *Configured Repositories*..

8 If needed, adjust the repository *Properties* as described in Section 9.3.2, “Managing Repository Properties” (page 197) or confirm your changes with *OK* to close the configuration dialog.

9 After having successfully added the repository for the add-on media, the software manager starts and you can install packages. Refer to Chapter 9, *Installing or Removing Software* (page 187) for details.

10.2 Binary Drivers

Some hardware needs binary-only drivers to function properly. If you have such hardware, refer to the release notes for more information about availability of binary drivers for your system. To read the release notes, open YaST and select *Miscellaneous > Release Notes*.

10.3 SUSE Software Development Kit (SDK) 11

SUSE Software Development Kit 11 is an add-on for SUSE Linux Enterprise 11. It is a complete tool kit for application development. In fact, to provide a comprehensive build system, SUSE Software Development Kit 11 includes all the open source tools that were used to build the SUSE Linux Enterprise Server product. It provides you - as a developer, independent software vendor (ISV), or independent hardware vendor (IHV) - with all the tools needed to port applications to all the platforms supported by SUSE Linux Enterprise Desktop and SUSE Linux Enterprise Server.

SUSE Software Development Kit also contains integrated development environments (IDEs), debuggers, code editors, and other related tools. It supports most major

programming languages, including C, C++, Java, and most scripting languages. For your convenience, SUSE Software Development Kit includes multiple Perl packages that are not included in SUSE Linux Enterprise.

For detailed information, refer to http://www.novell.com/developer/sle_sdk.html. Use the YaST add-on installer and package manager to install SUSE Software Development Kit 11.

Accessing the Internet

If you have chosen not to configure Internet access during the installation, you can perform this task at any time using YaST. How to configure your computer to access the Internet depends on your environment. If the computer you are installing is part of a network which already is connected to the Internet, the only thing to do is to link your machine to the network. If you are installing a machine that is directly connected to the Internet, the hardware and the access to the Internet Service Provider (ISP) needs to be set up.

Please refer to the checklists below to make sure you have all the necessary data ready before starting to configure the Internet access.

11.1 Direct Internet Connection

When your computer is directly connected to the Internet, you first need to configure the hardware that is used for this task. This can either be an internal device (such as an ISDN card) or an external device (for example, a modem). In most cases it is detected automatically.

Next, you need to enter the data provided by your ISP (such as login credentials, gateway, or name server, for example). You should have received a data sheet from your ISP where all the necessary data is listed.

If you have successfully configured your hardware and ISP data, use the `NetworkManager` for managing the internet connection. See Chapter 26, *Using NetworkManager* (↑*Administration Guide*) for details.

11.1.1 Checklist DSL

There are different types of DSL devices available that use different point-to-point protocol (PPP) methods:

- a regular ethernet card connected to the external DSL modem uses PPP over Ethernet (PPPoE). In Austria the Point-to-Point Tunneling Protocol (PPTP) is used. With PPTP the external modem also has a static IP address.
- an internal DSL modem uses PPP over ATM (PPPoATM)
- an internal ADSL Fritz Card uses CAPI for ADSL

The DSL configuration module already contains the data for major ISPs in some countries. If your ISP is not listed, you will need to know how name resolving (DNS) and IP allocation is handled (in most cases this data is received automatically when connecting). Regardless whether you choose an ISP from the list or add a custom provider, you need to enter at least your login and password.

For configuration details, refer to Section “DSL” (Chapter 21, *Basic Networking, Administration Guide*).

11.1.2 Checklist ISDN

In case your internal ISDN card is not detected automatically you will need to know the vendor and the name of the device.

NOTE: ISDN Modem or Terminal Adapter

If you are using an external ISDN modem or terminal adapter, refer to Section 11.1.3, “Checklist Modem” (page 211) instead.

In order to configure the ISDN device you will need the following data:

- ISDN Protocol (depends on your country)
- Area code and phone number.
- Interface type (SyncPPP or RawIP). If unsure, select SyncPPP, because RawIP is only used in connection with certain telephone systems.

- Local and remote IP addresses for the dial-in server and the gateway, in the case that you were given a static IP address from your provider.
- The ISDN configuration module already contains the data for major ISPs in some countries. If your ISP is not listed, you will need to know how name resolving (DNS) and IP allocation is handled (in most cases this data is received automatically when connecting). Regardless whether you chose an ISP from the list or added a custom provider, you need to enter at least your login and password.

For configuration details, refer to Section “ISDN” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

11.1.3 Checklist Modem

If your modem is not detected automatically, you will need to know whether it is connected to a serial port or to a USB port. Please note that not all USB modems and internal modems are supported by SUSE® Linux Enterprise Server.

The modem configuration module already contains the data for major ISPs in some countries. If your ISP is not listed, you will need to know its dial-in number and how name resolving (DNS) and IP allocation is handled (in most cases this data is received automatically when connecting). Regardless whether you chose an ISP from the list or added a custom provider, you need to enter at least your login and password.

For configuration details, refer to Section “Modem” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

11.1.4 Checklist Cable Modem

Accessing the Internet through the TV cable requires a cable modem. Such a modem is connected to the computer via ethernet cable. Therefore it is only necessary to configure your network card accordingly. For details, refer to Section “Cable Modem” (Chapter 21, *Basic Networking*, ↑*Administration Guide*).

11.2 Internet Connection Via Network

If your machine is part of a network which is already connected to the Internet, it is very easy to gain Internet access (just configure your network card and connect your machine to the existing network and you are done). This not only applies to large company networks, but to small home networks as well. Even if the machine you are installing is only connected to a router (e.g. a DSL router) it is already part of a network. It is irrelevant whether you are using a wireless network adapter or a wired one.

NOTE: Routing and Name Services

In the following it is assumed that the network is connected to the Internet and provides routing and name services. In case these services are provided by a router, make sure the router is configured correctly before setting up the client.

11.2.1 Checklist Network

If your network provides DHCP (Dynamic Host Configuration Protocol) check the appropriate check box when setting up the network card and you are done (all parameters needed will be provided by the DHCP server).

If DHCP is not available, ask your network administrator for the following details:

- Hostname
- Name server
- Gateway

For configuration details for wired network cards, refer to Section “Configuring the Network Card with YaST” (Chapter 21, *Basic Networking*, ↑*Administration Guide*), for wireless network cards see Section “Configuration with YaST” (Chapter 18, *Wireless LAN*, ↑*Administration Guide*).

Managing Users with YaST

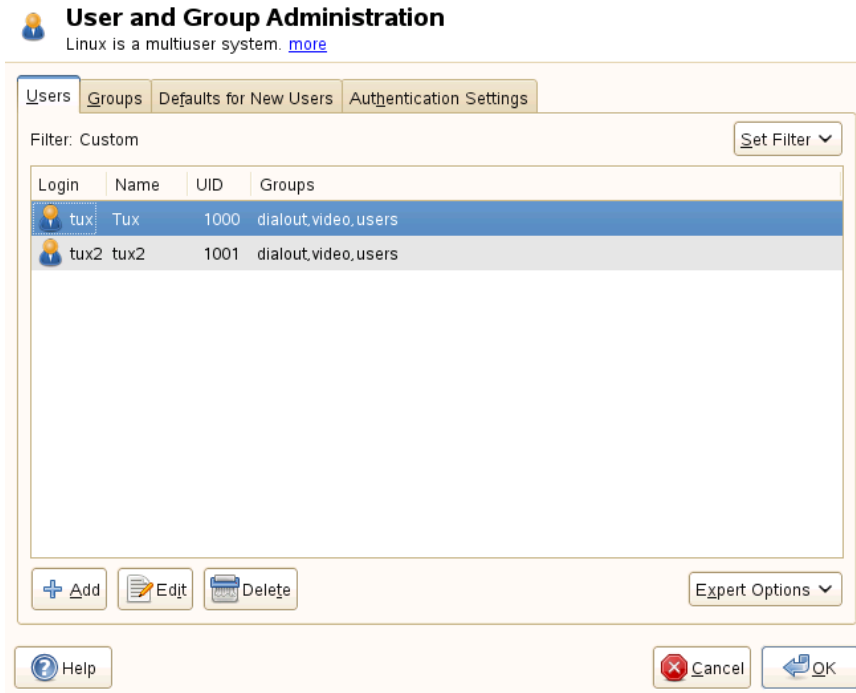
During installation, you chose a method for user authentication. This method is either local (via `/etc/passwd`) or, if a network connection is established, via NIS, LDAP, Kerberos or Samba (see Section 6.16.1.7, “User Authentication Method” (page 131). You can create or modify user accounts and change the authentication method with YaST at any time.

Every user is assigned a system-wide user ID (UID). Apart from the users which can log in to your machine, there are also a number of *system users* for internal use only. Each user is assigned to one or more groups. Similar to *system users*, there are also *system groups* for internal use.

12.1 User and Group Administration Dialog

To administer users or groups, start YaST and click *Security and Users > User and Group Management*. Alternatively, start the *User and Group Administration* dialog directly by running `yast2 users &` from a command line.

Figure 12.1: *YaST User and Group Administration*



Depending on the set of users you choose to view and modify with, the dialog (local users, network users, system users), the main window shows several tabs. These allow you to execute the following tasks:

Managing User Accounts

From the *Users* tab create, modify, delete or temporarily disable user accounts as described in Section 12.2, “Managing User Accounts” (page 215). Learn about advanced options like enforcing password policies, using encrypted home directories, using fingerprint authentication, or managing disk quotas in Section 12.3, “Additional Options for User Accounts” (page 217).

Changing Default Settings

Local users accounts are created according to the settings defined on the *Defaults for New Users* tab. Learn how to change the default group group assignment, or the default path and access permissions for home directories in Section 12.4, “Changing Default Settings for Local Users” (page 224).

Assigning Users to Groups

Learn how to change the group assignment for individual users in Section 12.5, “Assigning Users to Groups” (page 224).

Managing Groups

From the *Groups* tab, you can add, modify or delete existing groups. Refer to Section 12.6, “Managing Groups” (page 225) for information on how to do this.

Changing the User Authentication Method

When your machine is connected to a network that provides user authentication methods like NIS or LDAP, you can choose between several authentication methods on the *Authentication Settings* tab. For more information, refer to Section 12.7, “Changing the User Authentication Method” (page 226).

For user and group management, the dialog provides similar functionality. You can easily switch between the user and group administration view by choosing the appropriate tab at the top of the dialog.

Filter options allow you to define the set of users or groups you want to modify: On the *Users* or *Group* tab, click *Set Filter* to view and edit users or groups according to certain categories, such as *Local Users* or *LDAP Users*, for instance (if you are part of a network which uses LDAP). With *Set Filter > Customize Filter* you can also set up and use a custom filter.

Depending on the filter you choose, not all of the following options and functions will be available from the dialog.

12.2 Managing User Accounts

YaST offers to create, modify, delete or temporarily disable user accounts. Do not modify user accounts unless you are an experienced user or administrator.

NOTE: Changing User IDs of Existing Users

File ownership is bound to the user ID, not to the username. After a user ID change, the files in the user's home directory are automatically adjusted to reflect this change. However, after an ID change, the user no longer owns

the files he created elsewhere in the file system unless the file ownership for those files are manually modified.

In the following, learn how to set up default user accounts. For some further options, such as auto login, login without password, setting up encrypted home directories or managing quotas for users and groups, refer to Section 12.3, “Additional Options for User Accounts” (page 217).

Procedure 12.1: *Adding or Modifying User Accounts*

- 1 Open the YaST *User and Group Administration* dialog and click the *Users* tab.
- 2 With *Set Filter* define the set of users you want to manage. The dialog shows a list of users in the system and the groups the users belong to.
- 3 To modify options for an existing user, select an entry and click *Edit*.

To create a new user account, click *Add*.

- 4 Enter the appropriate user data on the first tab, such as *Username* (which is used for login) and *Password*. This data is sufficient to create a new user. If you click *OK* now, the system will automatically assign a user ID and set all other values according to the default.
- 5 If you want to adjust further details such as the user ID or the path to the user's home directory, do so on the *Details* tab.

If you need to relocate the home directory of an existing user, enter the path to the new home directory there and move the contents of the current home directory with *Move to New Location*. Otherwise, a new home directory is created without any of the existing data.

- 6 To force users to regularly change their password or set other password options, switch to *Password Settings* and adjust the options. For more details, refer to Section 12.3.2, “Enforcing Password Policies” (page 218).
- 7 If all options are set according to your wishes, click *OK*.
- 8 Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog. Click *OK* to close the administration

dialog and to save the changes. A newly added user can now log in to the system using the login name and password you created.

TIP: Matching User IDs

For a new (local) user on a laptop which also needs to integrate into a network environment where this user already has a user ID, it is useful to match the (local) user ID to the ID in the network. This ensures that the file ownership of the files the user creates “offline” is the same as if he had created them directly on the network.

Procedure 12.2: *Disabling or Deleting User Accounts*

- 1 Open the YaST *User and Group Administration* dialog and click the *Users* tab.
- 2 To temporarily disable a user account without deleting it, select the user from the list and click *Edit*. Activate *Disable User Login*. The user cannot log in to your machine until you enable the account again.
- 3 To delete a user account, select the user from the list and click *Delete*. Choose if you also want to delete the user's home directory or if you want to retain the data.

12.3 Additional Options for User Accounts

In addition to the settings for a default user account, SUSE® Linux Enterprise Server offers further options, such as options to enforce password policies, use encrypted home directories or define disk quotas for users and groups.

12.3.1 Automatic Login and Passwordless Login

If you use the KDE or GNOME desktop environment you can configure *Auto Login* for a certain user as well as *Passwordless Login* for all users. Auto login causes a user to become automatically logged in to the desktop environment on boot. This functionality can only be activated for one user at a time. Login without password

allows all users to log in to the system after they have entered their username in the login manager.

WARNING: Security Risk

Enabling *Auto Login* or *Passwordless Login* on a machine that can be accessed by more than one person is a security risk. Without the need to authenticate, any user can gain access to your system and your data. If your system contains confidential data, do not use this functionality.

If you want to activate auto login or login without password, access these functions in the YaST *User and Group Administration* with *Expert Options > Login Settings*.

12.3.2 Enforcing Password Policies

On any system with multiple users, it is a good idea to enforce at least basic password security policies. Users should change their passwords regularly and use strong passwords that cannot easily be exploited. For local users, proceed as follows:

Procedure 12.3: *Configuring Password Settings*

- 1 Open the YaST *User and Group Administration* dialog and select the *Users* tab.
- 2 Select the user for which to change the password options and click *Edit*.
- 3 Switch to the *Password Settings* tab. The user's last password change is displayed on the tab.
- 4 To make the user change his password at next login, activate *Force Password Change*.
- 5 To enforce password rotation, set a *Maximum Number of Days for the Same Password* and a *Minimum Number of Days for the Same Password*.
- 6 To remind the user to change his password before it expires, set a number of *Days before Password Expiration to Issue Warning*.
- 7 To restrict the period of time the user can log in after his password has expired, change the value in *Days after Password Expires with Usable Login*.

- 8 You can also specify a certain expiration date for a password. Enter the *Expiration Date* in *YYYY-MM-DD* format.
- 9 For more information about the options and about the default values, click *Help*.
- 10 Apply your changes with *OK*.

12.3.3 Managing Encrypted Home Directories

To protect data in home directories against theft and hard disk removal, you can create encrypted home directories for users. These are encrypted with LUKS (Linux Unified Key Setup), which results in an image and an image key being generated for the user. The image key is protected with the user's login password. When the user logs into the system, the encrypted home directory is mounted and the contents are made available to the user.

NOTE: Fingerprint Reader Devices and Encrypted Home Directories

If you want to use a fingerprint reader device, you must not use encrypted home directories. Otherwise logging in will fail, because decrypting during login is not possible in combination with an active fingerprint reader device.

With YaST, you can create encrypted home directories for new or existing users. To encrypt or modify encrypted home directories of already existing users, you need to know the user's current login password. By default, all existing user data is copied to the new encrypted home directory, but it is not deleted from the unencrypted directory.

WARNING: Security Restrictions

Encrypting a user's home directory does not provide strong security from other users. If strong security is required, the system should not be physically shared.


Find background information about encrypted home directories and which actions to take for stronger security in Section “Using Encrypted Home Directories” (Chapter 11, *Encrypting Partitions and Files*, ↑*Security Guide*).

Procedure 12.4: *Creating Encrypted Home Directories*

- 1 Open the YaST *User and Group Management* dialog and click the *Users* tab.
- 2 To encrypt the home directory of an existing user, select the user and click *Edit*.

Otherwise, click *Add* to create a new user account and enter the appropriate user data on the first tab.

- 3 In the *Details* tab, activate *Use Encrypted Home Directory*. With *Directory Size in MB*, specify the size of the encrypted image file to be created for this user.



The image shows the 'Existing Local User' dialog box in YaST, specifically the 'Details' tab. The dialog has four tabs: 'User Data', 'Details' (selected), 'Password Settings', and 'Plug-Ins'. The 'User ID (uid)' field contains '1000'. The 'Home Directory' field contains '/home/tux' with a 'Browse...' button. Below it, the 'Move to New Location' checkbox is checked. The 'Directory Size in MB' section has the 'Use Encrypted Home Directory' checkbox checked and a dropdown menu set to '100'. There is an 'Additional User Information' text field. The 'Login Shell' dropdown is set to '/bin/bash'. The 'Default Group' dropdown is set to 'users'. On the right, the 'Additional Groups' list includes checkboxes for 'users', 'at', 'audio', 'avahi', 'beagleindex', 'bin', 'cdrom', 'console', 'daemon', 'dialout' (checked), 'disk', 'floppy', 'ftp', 'games', and 'gdm'. At the bottom, there are 'Help', 'Cancel', and 'OK' buttons.

- 4 Apply your settings with *OK*.
- 5 Enter the user's current login password to proceed if YaST prompts for it.
- 6 Click *Expert Options > Write Changes Now* to save all changes without exiting the administration dialog. Click *OK* to close the administration dialog and save the changes.

Procedure 12.5: *Modifying or Disabling Encrypted Home Directories*

Of course, you can also disable the encryption of a home directory or change the size of the image file at any time.

- 1 Open the YaST *User and Group Administration* dialog in the *Users* view.
- 2 Select a user from the list and click *Edit*.
- 3 If you want to disable the encryption, switch to the *Details* tab and disable *Use Encrypted Home Directory*.

If you need to enlarge or reduce the size of the encrypted image file for this user, change the *Directory Size in MB*.
- 4 Apply your settings with *OK*.
- 5 Enter the user's current login password to proceed if YaST prompts for it.
- 6 Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog. Click *OK* to close the administration dialog and to save the changes.

12.3.4 Using Fingerprint Authentication

If your system includes a fingerprint reader you can use biometric authentication in addition to standard authentication via login and password. After registering their fingerprint, users can log in to the system either by swiping a finger on the fingerprint reader or by typing in a password.

Fingerprints can be registered with YaST. Find detailed information about configuration and use of fingerprint authentication in Chapter 7, *Using the Fingerprint Reader* (↑*Security Guide*). For a list of supported devices, refer to <http://www.freedesktop.org/wiki/Software/fprint/libfprint>.

12.3.5 Managing Quotas

To prevent system capacities from being exhausted without notification, system administrators can set up quotas for users or groups. Quotas can be defined for one

or more file systems and restrict the amount of disk space that can be used and the number of inodes (index nodes) that can be created there. Inodes are data structures on a file system that store basic information about a regular file, directory, or other file system object. They store all attributes of a file system object (like user and group ownership, read, write, or execute permissions), except file name and contents.

SUSE Linux Enterprise Server allows usage of `soft` and `hard` quotas. Soft quotas usually define a warning level at which users are informed that they are nearing their limit, whereas hard quotas define the limit at which write requests are denied. Additionally, grace intervals can be defined that allow users or groups to temporarily violate their quotas by certain amounts.

Procedure 12.6: *Enabling Quota Support for a Partition*

In order to configure quotas for certain users and groups, you need to enable quota support for the respective partition in the YaST Expert Partitioner first.

- 1** In YaST, select *System > Partitioner* and click *Yes* to proceed.
- 2** In the *Expert Partitioner*, select the partition for which to enable quotas and click *Edit*.
- 3** Click *Fstab Options* and activate *Enable Quota Support*. If the `quota` package is not already installed, it will be installed once you confirm the respective message with *Yes*.
- 4** Confirm your changes and leave the *Expert Partitioner*.

Procedure 12.7: *Setting Up Quotas for Users or Groups*

Now you can define soft or hard quotas for specific users or groups and set time periods as grace intervals.

- 1** In the YaST *User and Group Administration*, select the user or the group you want to set the quotas for and click *Edit*.
- 2** On the *Plug-Ins* tab, select the *Manage User Quota* entry and click *Launch* to open the *Quota Configuration* dialog.
- 3** From *File System*, select the partition to which the quota should apply.



Quota Configuration

Here, configure quota settings of the user on selected file systems. [more](#)

File System:

/dev/sda6

Size Limits

Soft limit:

5

Hard limit:

8

Days:

0

Hours:

0

Minutes:

0

Seconds:

0

I-nodes Limits

Soft limit:

2

Hard limit:

4

Days:

0

Hours:

0

Minutes:

0

Seconds:

0



- Below *Size Limits*, restrict the amount of disk space. Enter the number of 1 KB blocks the user or group may have on this partition. Specify a *Soft Limit* and a *Hard Limit* value.
- Additionally, you can restrict the number of inodes the user or group may have on the partition. Below *Inodes Limits*, enter a *Soft Limit* and *Hard Limit*.
- You can only define grace intervals if the user or group has already exceeded the soft limit specified for size or inodes. Otherwise, the time-related input fields are not activated. Specify the time period for which the user or group is allowed to exceed the limits set above.
- Confirm your settings with *OK*.
- Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog. Click *OK* to close the administration dialog and to save the changes.

SUSE Linux Enterprise Server also ships command line tools like `repquota` or `warnquota` with which system administrators can control the disk usage or send e-

mail notifications to users exceeding their quota. With `quota_nld`, administrators can also forward kernel messages about exceeded quotas to D-BUS. For more information, refer to the `repquota`, the `warnquota` and the `quota_nld` man page.

12.4 Changing Default Settings for Local Users

When creating new local users, several default settings are used by YaST. These include, for example, the primary group and the secondary groups the user belongs to, or the access permissions of the user's home directory. You can change these default settings to meet your requirements:

- 1 Open the *YaST User and Group Administration* dialog and select the *Defaults for New Users* tab.
- 2 To change the primary group the new users should automatically belong to, select another group from *Default Group*.
- 3 To modify the secondary groups for new users, add or change groups in *Secondary Groups*. The group names must be separated by commas.
- 4 If you do not want to use `/home/username` as default path for new users' home directories, modify the *Path Prefix for Home Directory*.
- 5 To change the default permission modes for newly created home directories, adjust the umask value in *Umask for Home Directory*. For more information about umask, refer to Chapter 10, *Access Control Lists in Linux* (↑*Security Guide*) and to the `umask` man page.
- 6 For information about the individual options, click *Help*.
- 7 Apply your changes with *OK*.

12.5 Assigning Users to Groups

Local users are assigned to several groups according to the default settings which you can access from the *User and Group Administration* dialog on the *Defaults for*

New Users tab. In the following, learn how to modify an individual user's group assignment. If you need to change the default group assignments for new users, refer to Section 12.4, “Changing Default Settings for Local Users” (page 224).

Procedure 12.8: *Changing a User's Group Assignment*


- 1 Open the YaST *User and Group Administration* dialog and click the *Users* tab. It shows a list of users and of the groups the users belong to.
- 2 Click *Edit* and switch to the *Details* tab.
- 3 To change the primary group the user belongs to, click *Default Group* and select the group from the list.
- 4 To assign the user additional secondary groups, activate the corresponding check boxes in the *Additional Groups* list.
- 5 Click *OK* to apply your changes.
- 6 Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog. Click *OK* to close the administration dialog and save the changes.

12.6 Managing Groups

With YaST you can also easily add, modify or delete groups.

Procedure 12.9: *Creating and Modifying Groups*

- 1 Open the YaST *User and Group Management* dialog and click the *Groups* tab.
- 2 With *Set Filter* define the set of groups you want to manage. The dialog shows a list of groups in the system.
- 3 To create a new group, click *Add*.
- 4 To modify an existing group, select the group and click *Edit*.
- 5 In the following dialog, enter or change the data. The list on the right shows an overview of all available users and system users which can be members of the group.

 **Existing Local Group**
Enter the group data here. [more](#)

Group Data

Plug-Ins

Group Name:

Group ID (gid):

Password:

Confirm Password:

Group Members:

☐ at
☐ avahi
☐ beagleindex
☐ bin
☐ daemon
☐ dnsmasq
☒ games
☒ tux
☒ tux2

Help

Cancel

OK

- 6 To add existing users to a new group select them from the list of possible *Group Members* by checking the corresponding box. To remove them from the group uncheck the box.
- 7 Click *OK* to apply your changes.
- 8 Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog.

In order to delete a group, it must not contain any group members. To delete a group, select it from the list and click *Delete*. Click *Expert Options > Write Changes Now* to save all changes without exiting the *User and Group Administration* dialog. Click *OK* to close the administration dialog and to save the changes.

12.7 Changing the User Authentication Method

When your machine is connected to a network, you can change the authentication method you set during installation. The following options are available:

NIS

Users are administered centrally on a NIS server for all systems in the network. For details, see Chapter 3, *Using NIS* (↑*Security Guide*).

LDAP

Users are administered centrally on an LDAP server for all systems in the network. For details about LDAP, see Chapter 4, *LDAP—A Directory Service* (↑*Security Guide*).

You can manage LDAP users with the YaST user module. All other LDAP settings, including the default settings for LDAP users, have to be defined with the YaST LDAP client module as described in Section “Configuring an LDAP Client with YaST” (Chapter 4, *LDAP—A Directory Service*, ↑*Security Guide*) .

Kerberos

With Kerberos, a user registers once and then is trusted in the entire network for the rest of the session.

Samba

SMB authentication is often used in mixed Linux and Windows networks. For details, see Chapter 27, *Samba* (↑*Administration Guide*).

eDirectory LDAP

eDirectory authentication is used in Novell networks.

To change the authentication method, proceed as follows:

- 1 Open the *User and Group Administration* dialog in YaST.
- 2 Click the *Authentication Settings* tab to show an overview of the available authentication methods and the current settings.
- 3 To change the authentication method, click *Configure* and select the authentication method you want to modify. This takes you directly to the client configuration modules in YaST. For information about the configuration of the appropriate client, refer to the following sections:

NIS: Section “Configuring NIS Clients” (Chapter 3, *Using NIS*, ↑*Security Guide*)

LDAP: Section “Configuring an LDAP Client with YaST” (Chapter 4, *LDAP—A Directory Service*, ↑*Security Guide*)

Samba: Section “Configuring a Samba Client with YaST” (Chapter 27, *Samba*,
↑*Administration Guide*)

- 4** After accepting the configuration, return to the *User and Group Administration* overview.
- 5** Click *OK* to close the administration dialog.

Changing Language and Country Settings with YaST

13

Working in different countries or having to work in a multilingual environment requires your computer to be set up to support this. SUSE® Linux Enterprise Server can handle different `locales` in parallel. A locale is a set of parameters that defines the language and country settings reflected in the user interface.

The main system language was selected during installation and keyboard and time zone settings were adjusted. However, you can install additional languages on your system and determine which of the installed languages should be the default.

For those tasks, use the YaST language module as described in Section 13.1, “Changing the System Language” (page 230). Install secondary languages to get optional localizations if you need to start applications or desktops in languages other than the primary one.

Apart from that, the YaST timezone module allows you to adjust your country and timezone settings accordingly. It also lets you synchronize your system clock against a time server. For details, refer to Section 13.2, “Changing the Country and Time Settings” (page 234).

13.1 Changing the System Language

Depending on how you use your desktop and whether you want to switch the entire system to another language or just the desktop environment itself, there are several ways to achieve this:

Changing the System Language Globally

Proceed as described in Section 13.1.1, “Modifying System Languages with YaST” (page 230) and Section 13.1.2, “Switching the Default System Language” (page 233) to install additional localized packages with YaST and to set the default language. Changes are effective after relogin. To ensure that the entire system reflects the change, reboot the system or close and restart all running services, applications, and programs.

Changing the Language for the Desktop Only

Provided you have previously installed the desired language packages for your desktop environment with YaST as described below, you can switch the language of your desktop using the desktop's control center. After the X server has been restarted, your entire desktop reflects your new choice of language. Applications not belonging to your desktop framework are not affected by this change and may still appear in the language that was set in YaST.

Temporarily Switching Languages for One Application Only

You can also run a single application in another language (that has already been installed with YaST). To do so, start it from the command line by specifying the language code as described in Section 13.1.3, “Switching Languages for Individual Applications” (page 233).

13.1.1 Modifying System Languages with YaST

YaST knows two different language categories:

Primary Language

The primary language set in YaST applies to the entire system, including YaST and the desktop environment. This language is used whenever available unless you manually specify another language.

Secondary Languages

Install secondary languages to make your system multilingual. Languages installed as secondary languages can be selected manually for a specific situation. For example, use a secondary language to start an application in a certain language in order to do word processing in this language.

Before installing additional languages, determine which of them should be the default system language (primary language) after you have installed them.

To access the YaST language module, start YaST and click *System > Language*. Alternatively, start the *Languages* dialog directly by running `yast2 language` & as user `root` from a command line.



Procedure 13.1: *Installing Additional Languages*

When installing additional languages, YaST also allows you to set different locale settings for the user `root`, see Step 4 (page 232). The option *Locale Settings*

for *User root* determines how the locale variables (LC_*) in the file `/etc/sysconfig/language` are set for *root*. You can either set them to the same locale as for normal users, keep it unaffected by any language changes or only set the variable `RC_LC_CTYPE` to the same values as for the normal users. This variable sets the localization for language-specific function calls.

- 1 To add additional languages in the YaST language module, select the *Secondary Languages* you wish to install.
- 2 To make a language the default language, set it as *Primary Language*.
- 3 Additionally, adapt the keyboard to the new primary language and adjust the time zone, if appropriate.

TIP

For advanced keyboard or time zone settings, select *Hardware > Keyboard Layout* or *System > Date and Time* in YaST to start the respective dialogs. For more information, refer to Section 13.2, “Changing the Country and Time Settings” (page 234).

- 4 To change language settings specific to the user *root*, click *Details*.
 - 4a Set *Locale Settings for User root* to the desired value. For more information, click *Help*.
 - 4b Decide if you want to *Use UTF-8 Encoding* for *root* or not.
- 5 If your locale was not included in the list of primary languages available, try specifying it with *Detailed Locale Setting*. However, some of these localizations may be incomplete.
- 6 Confirm your changes in the dialogs with *OK*. If you have selected secondary languages, YaST installs the localized software packages for the additional languages.

The system is now multilingual. However, to start an application in a language other than the primary one, you need to set the desired language explicitly as explained in Section 13.1.3, “Switching Languages for Individual Applications” (page 233).

13.1.2 Switching the Default System Language

- 1 To globally switch the default system language, start the YaST language module.
- 2 Select the desired new system language as *Primary Language*.

IMPORTANT: Deleting Former System Languages

If you switch to a different primary language, the localized software packages for the former primary language will be removed from the system. If you want to switch the default system language but want to keep the former primary language as additional language, add it as *Secondary Language* by enabling the respective check box.

- 3 Adjust the keyboard and time zone options as desired.
- 4 Confirm your changes with *OK*.
- 5 After YaST has applied the changes, restart any X sessions (for example, by logging out and logging in again) to make YaST and the desktop applications reflect your new language settings.

13.1.3 Switching Languages for Individual Applications

After you have installed the respective language with YaST, you can run a single application in another language.

Standard X and GNOME Applications

Start the application from the command line by using the following command:

```
LANG=language application
```

For example, to start f-spot in German, run `LANG=de_DE f-spot`. For other languages, use the appropriate language code. Get a list of all language codes available with the `locale -av` command.

KDE Applications

Start the application from the command line by using the following command:

`KDE_LANG=language application`

For example, to start digiKam in German, run `KDE_LANG=de digikam`. For other languages, use the appropriate language code.

13.2 Changing the Country and Time Settings

Using the YaST date and time module, adjust your system date, clock and time zone information to the area you are working in. To access the YaST module, start YaST and click *System > Date and Time*. Alternatively, start the *Clock and Time Zone* dialog directly by running `yast2 timezone` & as user `root` from a command line.



First, select a general region, such as *Europe*. Choose an appropriate country that matches the one you are working in, for example, *Germany*.

Depending on which operating systems run on your workstation, adjust the hardware clock settings accordingly:

- If you run another operating system on your machine, such as Microsoft Windows*, it is likely your system does not use UTC, but local time. In this case, uncheck *Hardware Clock Set To UTC*.
- If you only run Linux on your machine, set the hardware clock to UTC and have the switch from standard time to daylight saving time performed automatically.

You can change the date and time manually or opt for synchronizing your machine against an NTP server, either permanently or just for adjusting your hardware clock.

Procedure 13.2: *Manually Adjusting Time and Date*

- 1 In the YaST timezone module, click *Change* to set date and time.
- 2 Select *Manually* and enter date and time values.
- 3 Confirm your changes with *Accept*.

Procedure 13.3: *Setting Date and Time With NTP Server*

- 1 Click *Change* to set date and time.
- 2 Select *Synchronize with NTP Server*.
- 3 Enter the address of an NTP server, if not already populated.



Change Date and Time

The current system time and date are displayed. [more](#)

☐ Manually

Current Time:
 :

Current Date:
 - -

☒ Synchronize with NTP Server

NTP Server Address:

☒ Save NTP Configuration

- 4 Click *Synchronize Now*, to get your system time set correctly.
- 5 If you want to make use of NTP permanently, enable *Save NTP Configuration*.
- 6 With the *Configure* button, you can open the advanced NTP configuration. For details, see Section “Configuring an NTP Client with YaST” (Chapter 23, *Time Synchronization with NTP*, ↑*Administration Guide*).
- 7 Confirm your changes with *Accept*.

Remote Installation

SUSE® Linux Enterprise Server can be installed in different ways. As well as the usual media installation covered in Chapter 6, *Installation with YaST* (page 93), you can choose from various network-based approaches or even take a completely hands-off approach to the installation of SUSE Linux Enterprise Server.

Each method is introduced by means of two short checklists: one listing the prerequisites for this method and the other illustrating the basic procedure. More detail is then provided for all the techniques used in these installation scenarios.

NOTE

In the following sections, the system to hold your new SUSE Linux Enterprise Server installation is referred to as *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 and DVD, and network servers distributing the installation data in your network.

14.1 Installation 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 this scenario. If in need of detailed instructions for a particular step, follow the links provided for each one of them.

IMPORTANT

The configuration of the X Window System is not part of any remote installation process. After the installation has finished, log in to the target system as `root`, enter `telinit 3`, and start `SaX2` to configure the graphics hardware.

14.1.1 Simple Remote Installation via VNC —Static Network Configuration

This type of installation still requires some degree of physical access to the target system to boot for installation. The installation itself is entirely controlled by a remote workstation using VNC to connect to the installation program. User interaction is required as with the manual installation in Chapter 6, *Installation with YaST* (page 93).

For this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- Target system with working network connection.
- Controlling system with working network connection and VNC viewer software or Java-enabled browser (Firefox, Konqueror, Internet Explorer, Opera, etc.).
- Physical boot medium (CD, DVD, or USB flash drive) for booting the target system.
- Valid static IP addresses already assigned to the repository and the controlling system.
- Valid static IP address to assign to the target system.
- When installing over VNC, X11 is not configured and output is redirected to your local machine. To use `SaX2`, use `export DISPLAY=:0 && sax2 -a -r`

To perform this kind of installation, proceed as follows:

- 1 Set up the repository as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server. For an SMB repository, refer to Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2 Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- 3 When the boot screen of the target system appears, use the boot options prompt to set the appropriate VNC options and the address of the repository. This is described in detail in Section 14.4, “Booting the Target System for Installation” (page 266).

The target system boots to a text-based environment, giving the network address and display number under which the graphical installation environment can be addressed by any VNC viewer application or browser. VNC installations announce themselves over OpenSLP and if the firewall settings permit, they can be found using Konqueror in `service:/` or `slp:/` mode.

- 4 On the controlling workstation, open a VNC viewing application or Web browser and connect to the target system as described in Section 14.5.1, “VNC Installation” (page 270).
- 5 Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.
- 6 Finish the installation.

14.1.2 Simple Remote Installation via VNC —Dynamic Network Configuration

This type of installation still requires some degree of physical access to the target system to boot for installation. The network configuration is made with DHCP. The installation itself is entirely controlled from a remote workstation using VNC to connect to the installer, but still requires user interaction for the actual configuration efforts.

For this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- Target system with working network connection.
- Controlling system with working network connection and VNC viewer software or Java-enabled browser (Firefox, Konqueror, Internet Explorer, or Opera).
- Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- Running DHCP server providing IP addresses.
- When installing over VNC, X11 is not configured and output is redirected to your local machine. To use SaX2, use `export DISPLAY=:0 && sax2 -a -r`

To perform this kind of installation, proceed as follows:

- 1** Set up the repository as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server. For an SMB repository, refer to Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2** Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- 3** When the boot screen of the target system appears, use the boot options prompt to set the appropriate VNC options and the address of the repository. This is described in detail in Section 14.4, “Booting the Target System for Installation” (page 266).

The target system boots to a text-based environment, giving the network address and display number under which the graphical installation environment can be addressed by any VNC viewer application or browser. VNC installations announce themselves over OpenSLP and if the firewall settings permit, they can be found using Konqueror in `service:/` or `slp:/` mode.

- 4** On the controlling workstation, open a VNC viewing application or Web browser and connect to the target system as described in Section 14.5.1, “VNC Installation” (page 270).
- 5** Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.

6 Finish the installation.

14.1.3 Remote Installation via VNC—PXE Boot and Wake on LAN

This type of installation is completely hands-off. The target machine is started and booted remotely. User interaction is only needed for the actual installation. This approach is suitable for cross-site deployments.

To perform this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- TFTP server.
- Running DHCP server for your network.
- 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 VNC viewer software or Java-enabled browser (Firefox, Konqueror, Internet Explorer, or Opera).
- When installing over VNC, X11 is not configured and output is redirected to your local machine. To use SaX2, use `export DISPLAY=:0 && sax2 -a -r`

To perform this type of installation, proceed as follows:

- 1 Set up the repository as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server or configure an SMB repository as described in Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2 Set up a TFTP server to hold a boot image that can be pulled by the target system. This is described in Section 14.3.2, “Setting Up a TFTP Server” (page 258).
- 3 Set up a DHCP server to provide IP addresses to all machines and reveal the location of the TFTP server to the target system. This is described in Section 14.3.1, “Setting Up a DHCP Server” (page 255).

- 4 Prepare the target system for PXE boot. This is described in further detail in Section 14.3.5, “Preparing the Target System for PXE Boot” (page 265).
- 5 Initiate the boot process of the target system using Wake on LAN. This is described in Section 14.3.7, “Wake on LAN” (page 265).
- 6 On the controlling workstation, open a VNC viewing application or Web browser and connect to the target system as described in Section 14.5.1, “VNC Installation” (page 270).
- 7 Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.
- 8 Finish the installation.

14.1.4 Simple Remote Installation via SSH —Static Network Configuration

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 6, *Installation with YaST* (page 93).

For this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- Target system with working network connection.
- Controlling system with working network connection and working SSH client software.
- Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- Valid static IP addresses already assigned to the repository and the controlling system.
- Valid static IP address to assign to the target system.

- When installing over SSH, X11 is not configured and output is redirected to your local machine. To use SaX2, use `export DISPLAY=:0 && sax2 -a -r`

To perform this kind of installation, proceed as follows:

- 1 Set up the repository as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server. For an SMB repository, refer to Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2 Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- 3 When the boot screen of the target system appears, use the boot options prompt to set the appropriate parameters for network connection, address of the repository, and SSH enablement. This is described in detail in Section 14.4.2, “Using Custom Boot Options” (page 267).

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.

- 4 On the controlling workstation, open a terminal window and connect to the target system as described in Section 14.5.2.2, “Connecting to the Installation Program” (page 272).
- 5 Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.
- 6 Finish the installation.

14.1.5 Simple Remote Installation via SSH —Dynamic Network Configuration

This type of installation still requires some degree of physical access to the target system to boot for installation and 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, but still requires user interaction for the actual configuration efforts.

NOTE: Avoid Lost Connections After 2nd Step (Installation)

In the network settings dialog, check the *Traditional Method with ifup* and avoid NetworkManager. If not, your SSH connection will be lost during installation. Reset the settings to *User Controlled with NetworkManager* after your installation has finished.

For this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- Target system with working network connection.
- Controlling system with working network connection and working SSH client software.
- Physical boot medium (CD, DVD, or USB flash drive) for booting the target system.
- Running DHCP server providing IP addresses.
- When installing over SSH, X11 is not configured and output is redirected to your local machine. To use SaX2, use `export DISPLAY=:0 && sax2 -a -r`

To perform this kind of installation, proceed as follows:

- 1 Set up the repository source as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server. For an SMB repository, refer to Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2 Boot the target system using DVD1 of the SUSE Linux Enterprise Server media kit.
- 3 When the boot screen of the target system appears, use the boot options prompt to pass the appropriate parameters for network connection, location of the installation source, and SSH enablement. See Section 14.4.2, “Using Custom Boot Options” (page 267) for detailed instructions on the use of these parameters.

The target system boots to a text-based environment, giving you the network address under which the graphical installation environment can be addressed by any SSH client.

- 4 On the controlling workstation, open a terminal window and connect to the target system as described in Section 14.5.2.2, “Connecting to the Installation Program” (page 272).
- 5 Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.
- 6 Finish the installation.

14.1.6 Remote Installation via SSH—PXE Boot and Wake on LAN

This type of installation is completely hands-off. The target machine is started and booted remotely.

To perform this type of installation, make sure that the following requirements are met:

- Remote repository: NFS, HTTP, FTP, or SMB with working network connection.
- TFTP server.
- Running DHCP server for your network, providing a static IP to the host to install.
- 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 client software.
- When installing over SSH, X11 is not configured and output is redirected to your local machine. To use SaX2, use `export DISPLAY=:0 && sax2 -a -r`

To perform this type of installation, proceed as follows:

- 1 Set up the repository as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an NFS, HTTP, or FTP network server. For the configuration of an SMB repository, refer to Section 14.2.5, “Managing an SMB Repository” (page 253).
- 2 Set up a TFTP server to hold a boot image that can be pulled by the target system. This is described in Section 14.3.2, “Setting Up a TFTP Server” (page 258).

- 3 Set up a DHCP server to provide IP addresses to all machines and reveal the location of the TFTP server to the target system. This is described in Section 14.3.1, “Setting Up a DHCP Server” (page 255).
- 4 Prepare the target system for PXE boot. This is described in further detail in Section 14.3.5, “Preparing the Target System for PXE Boot” (page 265).
- 5 Initiate the boot process of the target system using Wake on LAN. This is described in Section 14.3.7, “Wake on LAN” (page 265).
- 6 On the controlling workstation, start an SSH client and connect to the target system as described in Section 14.5.2, “SSH Installation” (page 272).
- 7 Perform the installation as described in Chapter 6, *Installation with YaST* (page 93). Reconnect to the target system after it reboots for the final part of the installation.
- 8 Finish the installation.

14.2 Setting Up the Server Holding the Installation Sources

Depending on the operating system running on the machine to use 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 on SUSE Linux Enterprise Server 11 SP3 or openSUSE 11.1 and higher.

TIP

You can even use a Microsoft Windows machine as the installation server for your Linux deployment. See Section 14.2.5, “Managing an SMB Repository” (page 253) for details.

14.2.1 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 as `root` to the machine that should act as installation server.
- 2 Start *YaST > Miscellaneous > Installation Server*.
- 3 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.
- 4 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 Chapter 28, *Sharing File Systems with NFS* (↑*Administration Guide*).

TIP: Firewall Settings

Make sure that the firewall settings of your server system allow traffic on the ports for HTTP, NFS, and FTP. If they currently do not, enable *Open Port in Firewall* or check *Firewall Details* first.

- 5 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 that more add-on CDs or service pack CDs are required and should be added as extra repositories. To announce your installation server in the network via OpenSLP, activate the appropriate option.

TIP

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 just booted using the SLP boot option and find the network repository without any further configuration. For details on this option, refer to Section 14.4, “Booting the Target System for Installation” (page 266).

- 6 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 by hand if you have 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 and select *Add* in the overview of existing repositories to configure the new repository.

14.2.2 Setting Up an NFS Repository Manually

Setting up an NFS source for installation is basically done in two steps. In the first step, 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 later hold all installation data and change into this directory. For example:

```
mkdir install/product/productversion
cd install/product/productversion
```

Replace *product* with an abbreviation of the product name and *productversion* with a string that contains the product name and version.

- 3 For each DVD contained in the media kit execute the following commands:
 - 3a Copy the entire content of the installation DVD into the installation server directory:

```
cp -a /media/path_to_your_DVD_drive .
```

Replace *path_to_your_DVD_drive* with the actual path under which your DVD drive is addressed. Depending on the type of drive used in your system, this can be `cdrom`, `cdrecorder`, `dvd`, or `dvdrecorder`.

- 3b Rename the directory to the DVD number:

```
mv path_to_your_DVD_drive DVDx
```

Replace *x* with the actual number of your DVD.

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 hostnames of the machines to which to export the installation data. Instead of specifying hostnames here, you could also use wild cards, ranges of network addresses, or just 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.

If you prefer manually exporting the repository 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:

```
insserv /etc/init.d/nfsserver
```

- 4 Start the NFS server with `rcnfsserver start`. If you need to change the configuration of your NFS server later, modify the configuration file and restart the NFS daemon with `rcnfsserver restart`.

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 `rcslpd start`.

For more information about OpenSLP, refer to the package documentation located under `/usr/share/doc/packages/openslp/` or refer to Chapter 22, *SLP Services in the Network* (↑*Administration Guide*). More Information about NFS, refer to Chapter 28, *Sharing File Systems with NFS* (↑*Administration Guide*).

14.2.3 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 14.2.2, “Setting Up an NFS Repository Manually” (page 248).

2 Configure the FTP server to distribute the contents of your installation directory:

2a Log in as `root` and install the package `vsftpd` using the YaST software management.

2b Enter the FTP server root directory:

```
cd /srv/ftp
```

2c Create a subdirectory holding the installation sources in the FTP root directory:

```
mkdir repository
```

Replace `repository` with the product name.

2d 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`.

2e Start `vsftpd` with `vsftpd`.

3 Announce the repository via OpenSLP, if this is supported by your network setup:

3a 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 to the repository directory on your server. The `service:` line should be entered as one continuous line.

3b Start the OpenSLP daemon with `rcslpd start`.

TIP: Configuring an FTP Server with YaST

If you prefer using YaST over manually configuring the FTP installation server, refer to Chapter 31, *Setting up an FTP Server with YaST* (↑*Administration Guide*) for more information on how to use the YaST FTP server module.

14.2.4 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 14.2.2, “Setting Up an NFS Repository Manually” (page 248).
- 2** Configure the HTTP server to distribute the contents of your installation directory:
 - 2a** Install the Web server Apache as described in Section “Installation” (Chapter 30, *The Apache HTTP Server*, ↑*Administration Guide*).
 - 2b** Enter the root directory of the HTTP server (`/srv/www/htdocs`) and create the subdirectory that will hold the installation sources:

```
mkdir repository
```

Replace *repository* with the product name.

- 2c** 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
```

- 2d** 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
```

- 2e** Reload the HTTP server configuration using `rcapache2 reload`.

- 3** Announce the repository via OpenSLP, if this is supported by your network setup:

- 3a** 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.

- 3b** Start the OpenSLP daemon using `rcslpd restart`.

14.2.5 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 folder that will hold the entire installation tree and name it `INSTALL`, for example.
- 3 Export this share according the procedure outlined in your Windows documentation.
- 4 Enter this share and create a subfolder, called *product*. Replace *product* with the actual product name.
- 5 Enter the `INSTALL/product` folder and copy each DVD to a separate folder, such as `DVD1` and `DVD2`.

To use a 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), username, and password. Your syntax looks like this:

```
smb://workdomain;user:password@server/INSTALL/DVD1
```

After you hit **Enter**, YaST starts and you can perform the installation.

14.2.6 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 14.2.2, “Setting Up an NFS Repository Manually” (page 248), Section 14.2.3, “Setting Up an FTP Repository Manually” (page 251), or Section 14.2.4, “Setting Up an HTTP Repository Manually” (page 252).
- 4 Create subdirectories for each DVD.
- 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, *path_to_repository* with the source directory of your server, *product* with the product name, and *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 14.2.2, “Setting Up an NFS Repository Manually” (page 248), Section 14.2.3, “Setting Up an FTP Repository Manually” (page 251), or Section 14.2.4, “Setting Up an HTTP Repository Manually” (page 252).

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/product medium auto loop
```

14.3 Preparing the Boot of the Target System

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.

14.3.1 Setting Up a DHCP Server

There are two ways to set up a DHCP server. For SUSE Linux Enterprise Server, YaST provides a graphical interface to the process. Users can also manually edit the

configuration files. For more information about DHCP servers, see also Chapter 25, *DHCP* (↑*Administration Guide*).

14.3.1.1 Setting Up a DHCP Server with YaST

To announce the TFTP server's location to the network clients and specify the boot image file the installation target should use, add two declarations to your DHCP server configuration.

- 1 Log in as `root` to the machine hosting the DHCP server.
- 2 Start *YaST* > *Network Services* > *DHCP Server*.
- 3 Complete the setup wizard for basic DHCP server setup.
- 4 Select *Expert Settings* and select *Yes* when warned about leaving the start-up dialog.
- 5 In the *Configured Declarations* dialog, select the subnet in which the new system should be located and click *Edit*.
- 6 In the *Subnet Configuration* dialog select *Add* to add a new option to the subnet's configuration.
- 7 Select `filename` and enter `pxelinux.0` as the value.
- 8 Add another option (`next-server`) and set its value to the address of the TFTP server.
- 9 Select *OK* and *Finish* to complete the DHCP server configuration.

To configure DHCP to provide a static IP address to a specific host, enter the *Expert Settings* of the DHCP server configuration module (Step 4 (page 256)) and add a new declaration of the host type. Add the options `hardware` and `fixed-address` to this host declaration and provide the appropriate values.

14.3.1.2 Setting Up a DHCP Server Manually

All the DHCP server needs to do, apart from providing automatic address allocation to your network clients, is to announce the IP address of the TFTP server and the file that needs to be pulled in by the installation routines on the target machine.

- 1 Log in as `root` to the machine hosting the DHCP server.
- 2 Append the following lines to a subnet configuration of your DHCP server's configuration file located under `/etc/dhcpd.conf`:

```
subnet 192.168.1.0 netmask 255.255.255.0 {
    range dynamic-bootp 192.168.1.200 192.168.1.228;
    # PXE related stuff
    #
    # "next-server" defines the TFTP server that will be used
    next-server ip_tftp_server:
    #
    # "filename" specifies the pxelinux image on the TFTP server
    # the server runs in chroot under /srv/tftpboot
    filename "pxelinux.0";
}
```

Replace *ip_of_the_tftp_server* with the actual IP address of the TFTP server. For more information about the options available in `dhcpd.conf`, refer to the `dhcpd.conf` manual page.

- 3 Restart the DHCP server by executing `rcdhcpd restart`.

If you plan on using SSH for the remote control of a PXE and Wake on LAN installation, explicitly specify the IP address DHCP should provide to the installation target. To achieve this, modify the above mentioned DHCP configuration according to the following example:

```
group {
    # PXE related stuff
    #
    # "next-server" defines the TFTP server that will be used
    next-server ip_tftp_server:
    #
    # "filename" specifies the pxelinux image on the TFTP server
    # the server runs in chroot under /srv/tftpboot
    filename "pxelinux.0";
    host test {
        hardware ethernet mac_address;
        fixed-address some_ip_address;
    }
}
```

The host statement introduces the hostname of the installation target. To bind the hostname and IP address to a specific host, you must know and specify the system's hardware (MAC) address. Replace all the variables used in this example with the actual values that match your environment.

After restarting the DHCP server, it provides a static IP to the host specified, enabling you to connect to the system via SSH.

14.3.2 Setting Up a TFTP Server

Set up a TFTP server with YaST on SUSE Linux Enterprise Server and SUSE Linux Enterprise Server or set it up manually on any other Linux operating system that supports xinetd and TFTP. The TFTP server delivers the boot image to the target system once it boots and sends a request for it.

14.3.2.1 Setting Up a TFTP Server Using YaST

- 1 Log in as `root`.
- 2 Start *YaST > Network Services > TFTP Server* and install the requested package.
- 3 Click *Enable* to make sure that the server is started and included in the boot routines. No further action from your side is required to secure this. `xinetd` starts `tftpd` at boot time.
- 4 Click *Open Port in Firewall* to open the appropriate port in the firewall running on your machine. If there is no firewall running on your server, this option is not available.
- 5 Click *Browse* to browse for the boot image directory. The default directory `/tftpboot` is created and selected automatically.
- 6 Click *Finish* to apply your settings and start the server.

14.3.2.2 Setting Up a TFTP Server Manually

- 1 Log in as `root` and install the packages `tftp` and `xinetd`.
- 2 If unavailable, create `/srv/tftpboot` and `/srv/tftpboot/pxelinux.cfg` directories.
- 3 Add the appropriate files needed for the boot image as described in Section 14.3.3, “Using PXE Boot” (page 259).

- 4** Modify the configuration of `xinetd` located under `/etc/xinetd.d` to make sure that the TFTP server is started on boot:

4a If it does not exist, create a file called `tftp` under this directory with `touch tftp`. Then run `chmod 755 tftp`.

4b Open the file `tftp` and add the following lines:

```
service tftp
{
    socket_type          = dgram
    protocol             = udp
    wait                 = yes
    user                 = root
    server               = /usr/sbin/in.tftpd
    server_args          = -s /srv/tftpboot
    disable              = no
}
```

4c Save the file and restart `xinetd` with `rcxinetd restart`.

14.3.3 Using PXE Boot

Some technical background information as well as PXE's complete specifications are available in the Preboot Execution Environment (PXE) Specification (<http://www.pix.net/software/pxeboot/archive/pxespec.pdf>).

- 1** Change to the directory `boot/<architecture>/loader` of your installation repository and copy the `linux`, `initrd`, `message`, `biostest`, and `memtest` files to the `/srv/tftpboot` directory by entering the following:

```
cp -a linux initrd message biostest memtest /srv/tftpboot
```

- 2** Install the `syslinux` package directly from your installation DVDs with YaST.

- 3** Copy the `/usr/share/syslinux/pxelinux.0` file to the `/srv/tftpboot` directory by entering the following:

```
cp -a /usr/share/syslinux/pxelinux.0 /srv/tftpboot
```

- 4** Change to the directory of your installation repository and copy the `isolinux.cfg` file to `/srv/tftpboot/pxelinux.cfg/default` by entering the following:

```
cp -a boot/<architecture>/loader/isolinux.cfg /srv/tftpboot/pxelinux.cfg/  
default
```

- 5 Edit the `/srv/tftpboot/pxelinux.cfg/default` file and remove the lines beginning with `readinfo` and `framebuffer`.
- 6 Insert the following entries in the append lines of the default `failsafe` and `apic` labels:

```
insmod=kernel module
```

By means of this entry, enter the network Kernel module needed to support network installation on the PXE client. Replace *kernel module* with the appropriate module name for your network device.

```
netdevice=interface
```

This entry defines the client's network interface that must be used for the network installation. It is only necessary if the client is equipped with several network cards and must be adapted accordingly. In case of a single network card, this entry can be omitted.

```
install=nfs://ip_instserver/path_to_repository/DVD1
```

This entry defines the NFS server and the repository for the client installation. Replace *ip_instserver* with the actual IP address of your installation server. *path_to_repository* should be replaced with the actual path to the repository. HTTP, FTP, or SMB repositories are addressed in a similar manner, except for the protocol prefix, which should read `http`, `ftp`, or `smb`.

IMPORTANT

If you need to pass other boot options to the installation routines, such as SSH or VNC boot parameters, append them to the `install` entry. An overview of parameters and some examples are given in Section 14.4, “Booting the Target System for Installation” (page 266).

TIP: Changing Kernel and initrd Filenames

It is possible to use different filenames for Kernel and initrd images. This is useful if you want to provide different operating systems from the

same boot server. However, you should be aware that only one dot is permitted in the filenames that are provided by TFTP for the PXE boot.

An example `/srv/tftpbboot/pxelinux.cfg/default` file follows. Adjust the protocol prefix for the repository to match your network setup and specify your preferred method of connecting to the installer by adding the `vnc` and `vncpassword` or the `usessh` and `sshpasword` options to the `install` entry. The lines separated by `\` must be entered as one continuous line without a line break and without the `\`.

```
default hard disk

# default
label linux
    kernel linux
    append initrd=initrd ramdisk_size=65536 \
        install=nfs://ip_instserver/path_to_repository/product/DVD1

# repair
label repair
    kernel linux
    append initrd=initrd splash=silent repair=1 showopts

# rescue
label rescue
    kernel linux
    append initrd=initrd ramdisk_size=65536 rescue=1

# bios test
label firmware
    kernel linux
    append initrd=biostest,initrd splash=silent install=exec:/bin/
run_biostest showopts

# memory test
label memtest
    kernel memtest

# hard disk
label hard disk
    localboot 0

implicit      0
display       message
prompt        1
timeout       100
```

- 7 Replace `ip_instserver` and `path_to_repository` with the values used in your setup.

The following section serves as a short reference to the PXELINUX options used in this setup. Find more information about the options available in the documentation of the `syslinux` package located under `/usr/share/doc/packages/syslinux/`.

14.3.4 PXELINUX Configuration Options

The options listed here are a subset of all the options available for the PXELINUX configuration file.

`APPEND options...`

Add 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 -`

Append 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. If PXELINUX boots automatically, it acts as if the entries after `DEFAULT` had been typed in at the boot prompt, except the `auto` option is automatically added, indicating an automatic boot.

If no configuration file is present or no `DEFAULT` entry is present 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 14.1, “Generated and Added Kernel Command Line Options from `IFAPPEND`” (page 262):

Table 14.1: *Generated and Added Kernel Command Line Options from `IFAPPEND`*

Argument	Generated Kernel Command Line / Description
1	<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>
2	<p><code>BOOTIF=MAC_ADDRESS_OF_BOOT_INTERFACE</code></p> <p>This option is useful if you want to avoid timeouts when the installation server probes one LAN interface after the other until it gets a reply from a DHCP server. Using 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>
4	<p><code>SYSUUID=SYSTEM_UUID</code></p> <p>Adds UUIDs in lowercase hexadecimals, see <code>/usr/share/doc/packages/syslinux/pxelinux.txt</code></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 instead of 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.

Note that GRUB uses the following syntax:

```
title mytitle
  kernel my_kernelmy_kernel_options
  initrd myinitrd
```

PXELINUX uses the following syntax:

```
label mylabel
  kernel mykernel
  append myoptions
```

Labels are mangled as if they were filenames 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 filename.

The Kernel does not have to be a Linux Kernel; it can 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
0	Perform a normal boot
4	Perform a local boot with the Universal Network Driver Interface (UNDI) driver still resident in memory
5	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 as soon as 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 filenames to F11 and F12.

14.3.5 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 option in the BIOS. Otherwise this system would try to re-install itself every time you boot it.

14.3.6 Preparing the Target System for Wake on LAN

Wake on LAN (WOL) requires the appropriate BIOS option to be enabled prior to the installation. Also, note down the MAC address of the target system. This data is needed to initiate Wake on LAN.

14.3.7 Wake on LAN

Wake on LAN allows a machine to be turned on by a special network packet containing the machine's MAC address. Because every machine in the world has a unique MAC identifier, you do not need to worry about accidentally turning on the wrong machine.

IMPORTANT: Wake on LAN across Different Network Segments

If the controlling machine is not located in the same network segment as the installation target that should be awakened, either configure the WOL requests to be sent as multicasts or remotely control a machine on that network segment to act as the sender of these requests.

Users of SUSE Linux Enterprise Server can use a YaST module called WOL to easily configure Wake on LAN. Users of other versions of SUSE Linux-based operating systems can use a command line tool.

14.3.8 Wake on LAN with YaST

- 1 Log in as `root`.
- 2 Start *YaST > Network Services > WOL*.
- 3 Click *Add* and enter the hostname and MAC address of the target system.
- 4 To turn on this machine, select the appropriate entry and click *Wake up*.

14.4 Booting the Target System for Installation

Basically, there are two different ways to customize the boot process for installation apart from those mentioned under Section 14.3.7, “Wake on LAN” (page 265) and Section 14.3.3, “Using PXE Boot” (page 259). You can either use the default boot options and function keys or use the boot options prompt of the installation boot screen to pass any boot options that the installation Kernel might need on this particular hardware.

14.4.1 Using the Default Boot Options

The boot options are described in detail in Chapter 6, *Installation with YaST* (page 93). Generally, just 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 Section “Installation Problems” (Chapter 35, *Common Problems and Their Solutions*, ↑*Administration Guide*).

The menu bar at the bottom screen offers some advanced functionality needed in some setups. Using the F keys, you can specify additional options to pass to the installation routines without having to know the detailed syntax of these parameters (see Section 14.4.2, “Using Custom Boot Options” (page 267)). A detailed description of the available function keys is available at Section 6.6, “The Boot Screen on Machines Equipped with Traditional BIOS” (page 98).

14.4.2 Using Custom Boot Options

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

The following table lists all installation scenarios mentioned in this chapter with the required parameters for booting and the corresponding boot options. Just append all of them in the order they appear in this table to get one boot option string that is handed to the installation routines. For example (all in one line):

```
install=xxx netdevice=xxx hostip=xxx netmask=xxx vnc=xxx vncpassword=xxx
```

Replace all the values `xxx` in this string with the values appropriate for your setup.

Table 14.2: *Installation (Boot) Scenarios Used in This Chapter*

Installation Scenario	Parameters Needed for Booting	Boot Options
Chapter 6, <i>Installation with YaST</i> (page 93)	None: system boots automatically	None needed
Section 14.1.1, “Simple Remote Installation via VNC—Static Network Configuration” (page 238)	<ul style="list-style-type: none"> • Location of the installation server • Network device • IP address • Netmask • Gateway 	<ul style="list-style-type: none"> • <code>install=(nfs,http,ftp,smb)://path_to_instmedia</code> • <code>netdevice=some_netdevice</code> (only needed if several network

Installation Scenario	Parameters Needed for Booting	Boot Options
	<ul style="list-style-type: none"> • VNC enablement • VNC password 	<p>devices are available)</p> <ul style="list-style-type: none"> • <code>hostip=some_ip</code> • <code>netmask=some_netmask</code> • <code>gateway=ip_gateway</code> • <code>vnc=1</code> • <code>vncpassword=some_password</code>
Section 14.1.2, “Simple Remote Installation via VNC—Dynamic Network Configuration” (page 239)	<ul style="list-style-type: none"> • Location of the installation server • VNC enablement • VNC password 	<ul style="list-style-type: none"> • <code>install=(nfs,http,ftp,smb)://path_to_instmedia</code> • <code>vnc=1</code> • <code>vncpassword=some_password</code>
Section 14.1.3, “Remote Installation via VNC—PXE Boot and Wake on LAN” (page 241)	<ul style="list-style-type: none"> • Location of the installation server • Location of the TFTP server • VNC enablement • VNC password 	Not applicable; process managed through PXE and DHCP
Section 14.1.4, “Simple Remote Installation via SSH—Static Network Configuration” (page 242)	<ul style="list-style-type: none"> • Location of the installation server • Network device • IP address • Netmask • Gateway • SSH enablement • SSH password 	<ul style="list-style-type: none"> • <code>install=(nfs,http,ftp,smb)://path_to_instmedia</code> • <code>netdevice=some_netdevice</code> (only needed if several network devices are available) • <code>hostip=some_ip</code> • <code>netmask=some_netmask</code> • <code>gateway=ip_gateway</code> • <code>usessh=1</code> • <code>sshpassword=some_password</code>

Installation Scenario	Parameters Needed for Booting	Boot Options
Section 14.1.5, “Simple Remote Installation via SSH —Dynamic Network Configuration” (page 243)	<ul style="list-style-type: none"> • Location of the installation server • SSH enablement • SSH password 	<ul style="list-style-type: none"> • <code>install=(nfs,http,ftp,smb):// path_to_instmedia</code> • <code>usessh=1</code> • <code>sshpassword=some_password</code>
Section 14.1.6, “Remote Installation via SSH—PXE Boot and Wake on LAN” (page 245)	<ul style="list-style-type: none"> • Location of the installation server • Location of the TFTP server • SSH enablement • SSH password 	Not applicable; process managed through PXE and DHCP

TIP: More Information about linuxrc Boot Options

Find more information about the linuxrc boot options used for booting a Linux system at <http://en.opensuse.org/SDB:Linuxrc>.

14.4.2.1 Installing Add-On Products and Driver Updates

SUSE Linux Enterprise Server supports the installation of Add-On products providing extensions (for example the SUSE Linux Enterprise High Availability Extension), third-party products and drivers or additional software. In order to automatically install an Add-On product when deploying SUSE Linux Enterprise Server remotely, specify the `addon=REPOSITORY` parameter.

REPOSITORY needs to be a hosted repository that can be read by YaST (YaST2 or YUM (rpm-md)). ISO images are currently not supported.

TIP: Driver Updates

Driver Updates can be found at <http://drivers.suse.com/>. Not all driver updates are provided as repositories—some are only available as iso images and therefore cannot be installed with the `addon` parameter.

Instructions on how to install driver updates via iso image are available at http://drivers.suse.com/doc/kit_usage.html.

14.5 Monitoring the Installation Process

There are several options for remotely monitoring the installation process. If the proper boot options have been specified while booting for installation, either VNC or SSH can be used to control the installation and system configuration from a remote workstation.

14.5.1 VNC Installation

Using any VNC viewer software, 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 application or a Web browser.

14.5.1.1 Preparing for VNC Installation

All you need to do on the installation target to prepare for a VNC installation is to provide the appropriate boot options at the initial boot for installation (see Section 14.4.2, “Using Custom Boot Options” (page 267)). 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 without the need for any physical contact to the installation itself, provided your network setup and all machines support OpenSLP:

- 1 Start the KDE file and Web browser Konqueror.

- 2 Enter `service://yast.installation.suse` in the location bar. The target system then appears as an icon in the Konqueror screen. Clicking this icon launches the KDE VNC viewer in which to perform the installation. Alternatively, run your VNC viewer software with the IP address provided and add `:1` at the end of the IP address for the display the installation is running on.

14.5.1.2 Connecting to the Installation Program

Basically, there are two ways to connect to a VNC server (the installation target in this case). You can either start an independent VNC viewer application on any operating system or connect using a Java-enabled Web browser.

Using VNC, you can control the installation of a Linux system from any other operating system, including other Linux flavors, Windows, or Mac OS.

On a Linux machine, make sure that the package `tightvnc` is installed. On a Windows machine, install the Windows port of this application, which can be obtained at the TightVNC home page (<http://www.tightvnc.com/download.html>).

To connect to the installation program 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 as provided by the SLP browser or the installation program itself:

ip_address:display_number

A window opens on your desktop displaying the YaST screens as in a normal local installation.

Using a Web browser to connect to the installation program makes you totally independent of any VNC software or the underlying operating system. As long as the browser application has Java support enabled, you can use any browser (Firefox, Internet Explorer, Konqueror, Opera, etc.) to perform the installation of your Linux system.

To perform a VNC installation, proceed as follows:

- 1 Launch your preferred Web browser.

- 2 Enter the following at the address prompt:

```
http://ip_address_of_target:5801
```

- 3 Enter your VNC password when prompted to do so. The browser window now displays the YaST screens as in a normal local installation.

14.5.2 SSH Installation

Using SSH, you can remotely control the installation of your Linux machine using any SSH client software.

14.5.2.1 Preparing for SSH Installation

Apart from installing the appropriate software package (OpenSSH for Linux and PuTTY for Windows), you just need to pass the appropriate boot options to enable SSH for installation. See Section 14.4.2, “Using Custom Boot Options” (page 267) for details. OpenSSH is installed by default on any SUSE Linux-based operating system.

14.5.2.2 Connecting to the Installation Program

- 1 Retrieve the installation target's IP address. If you have physical access to the target machine, just take the IP address the installation routine provides at the console after the initial boot. Otherwise take the IP address that has been assigned to this particular host in the DHCP server configuration.
- 2 At a command line, enter the following command:

```
ssh -X root@  
ip_address_of_target
```

Replace *ip_address_of_target* with the actual IP address of the installation target.

- 3 When prompted for a username, enter `root`.
- 4 When prompted for the password, enter the password that has been set with the SSH boot option. After you have successfully authenticated, a command line prompt for the installation target appears.

- 5 Enter `yast` to launch the installation program. A window opens showing the normal YaST screens as described in Chapter 6, *Installation with YaST* (page 93).

Advanced Disk Setup

Sophisticated system configurations require specific disk setups. All common partitioning tasks can be done with YaST. 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 Chapter 7, *Managing Multipath I/O for Devices* (↑Storage Administration Guide) for details), and there is also the option to use iSCSI as a networked disk (read more about iSCSI in Chapter 14, *Mass Storage over IP Networks: iSCSI* (↑Storage Administration Guide)).

15.1 Using the YaST Partitioner


With the expert partitioner, shown in Figure 15.1, “The YaST Partitioner” (page 276), manually modify the partitioning of one or several hard disks. 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 the data loss is very high. Try to avoid

repartitioning your installed system and always do a complete backup of your data before attempting to do so.

Figure 15.1: *The YaST Partitioner*



TIP: IBM System z: Device Names

IBM System z recognize only DASD 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 lefthand *System View*. Use these views to gather information about existing storage configurations, or to configure functions like RAID, Volume Management, Crypt Files, or 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 starting from the bottom toward the top of the list (starting from the last partition of a hard disk toward the first). For example,

if you have three partitions, you cannot use the second exclusively for SUSE Linux Enterprise Server and retain the third and first for other operating systems.

15.1.1 Partition Types

TIP: IBM System z: Hard Disks

On the IBM System z platforms, SUSE Linux Enterprise Server supports SCSI hard disks as well as DASDs (direct access storage devices). While SCSI disks can be partitioned as described below, DASDs can have no more than three partition entries in their partition tables.

Every hard disk has a partition table with space for four entries. Every entry in the partition table corresponds to a primary partition or an extended partition. Only one extended partition entry is allowed, however.

A primary partition simply consists of a continuous range of cylinders (physical disk areas) assigned to a particular operating system. With primary partitions you would be limited to four partitions per hard disk, because more do not fit in the partition table. This is why extended partitions are used. Extended partitions are also continuous ranges of disk cylinders, but an extended partition may be divided into *logical partitions* itself. Logical partitions do not require entries in the partition table. In other words, an extended partition is a container for logical partitions.

If you need more than four partitions, create an extended partition as the fourth partition (or earlier). This extended partition should occupy the entire remaining free cylinder range. Then create multiple logical partitions within the extended partition. The maximum number of logical partitions is 63, independent of the disk type. It does not matter which types of partitions are used for Linux. Primary and logical partitions both function normally.

TIP: GPT Partition Table

If you need to create more than 4 primary partitions on one hard disk, you have to use the GPT partition type. This type removes the primary partitions number restriction, and supports partitions bigger than 2 TB as well.

To use GPT, run the YaST Partitioner, click the relevant disk name in the *System View* and choose *Expert > Create New Partition Table > GPT*.

15.1.2 Creating a Partition

To create a partition from scratch select *Hard Disks* and then a hard disk with free space. The actual modification can be done in the *Partitions* tab:

- 1 Select *Add* and specify the partition type (primary or extended). Create up to four primary partitions or up to three primary partitions and one extended partition. Within the extended partition, create several logical partitions (see Section 15.1.1, “Partition Types” (page 277)).
- 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*. For more information on supported file systems, see *Storage Administration Guide*.
- 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 15.1.3, “Editing a Partition” (page 281).
- 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.

15.1.2.1 Btrfs Partitioning

If you want to use Btrfs (see Chapter 4, *Snapshots/Rollback with Snapper* (↑*Administration Guide*) and *Storage Administration Guide*, chapter *Overview of File Systems in Linux* for more information on Btrfs) as your default file system for a newly installed system, click *Partitioning* on the *Installation Settings* screen, and check *Use Btrfs as Default Filesystem*. The installation system then suggests creating the `/boot` partition formatted with Ext3 file system, and the root `/` partition formatted with Btrfs holding a default set of subvolumes, which you can modify with the *Expert Partitioner* tool later.

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.

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` snapshots the *root* file system by default, it is reasonable to exclude specific directories from being snapshot, depending on the nature of data they hold. And that is why YaST suggests creating the following separate subvolumes.

Suggested Btrfs Subvolumes

`/tmp /var/tmp /var/run`

Directories with frequently changed content.

`/var/spool`

Contains user data, such as mails.

`/var/log`

Contains system and applications' log files which should never be rolled back.

`/var/crash`

Contains memory dumps of crashed kernels.

`/srv`

Contains data files belonging to FTP and HTTP servers.

`/opt`

Contains third party software.

TIP: Size of Btrfs Partition

Because saved snapshots require more disk space, it is recommended to reserve more space for Btrfs partition than for a partition not capable of snapshotting (such as Ext3). Recommended size for a root Btrfs partition with suggested subvolumes is 20GB.


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 remove existing subvolumes.

Procedure 15.1: *Btrfs Subvolumes with YaST*

- 1 Start the YaST *Expert Partitioner* with *System > Partitioner*.
- 2 Choose *BTRFS* in the left *System View* pane.
- 3 Select the Btrfs partition whose subvolumes you need to manage and click *Edit*.
- 4 Click *Subvolume Handling*. You can see a list off all existing subvolumes of the selected Btrfs partition. You can notice a number of `@/.snapshots/xyz/snapshot` entries — each of these subvolumes belongs to one existing snapshot.
- 5 Depending on whether you want to add or remove subvolumes, do the following:
 - 5a To remove a subvolume, select it from the list of *Existing Subvolumes* and click *Remove*.
 - 5b To add a new subvolume, enter its name to the *New Subvolume* text field and click *Add new*.

Figure 15.2: *Btrfs Subvolumes in YaST Partitioner*



6 Confirm with *OK* and *Finish*.

7 Leave the partitioner with *Finish*.

15.1.3 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 because of its advanced features. It supports copy-on-write functionality, creating snapshots, multi-device spanning, subvolumes, and other useful techniques. ReiserFS, JFS, XFS, and Ext3 are journaling file systems. These file systems are able to restore the system very quickly after a system crash, utilizing 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.

NOTE: Support for Ext4 Filesystem

Because Btrfs proved to be more efficient and scalable than Ext4, SUSE Linux Enterprise Server SP2 supports only read-only access to Ext4 partitions. It is, however, still possible to access Ext4 partitions in a read-write mode — you need to install the

`ext4-writeable` package. Please note that this operation is not supported and taints the Kernel.

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 more memory to your system instead of adding more 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 Chapter 11, *Encrypting Partitions and Files* (↑*Security Guide*).

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 System z: Mounting by path

Since mounting by ID causes problems on IBM System z when using disk-to-disk copying for cloning purposes, devices are mounted by path in `/etc/fstab` on IBM System 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 Section 12.3.5, “Managing Quotas” (page 221).

3 Select *Finish* to save the changes.

NOTE: Resize Filesystems

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.

15.1.4 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 removes all the 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.

15.1.5 Advanced Options

After you select the hostname 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 have 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.

15.1.6 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.

TIP: 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.

15.1.6.1 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 killing 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, just add this partition with YaST. If you do not have a partition available, you may also just 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 15.2: *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, if you want to add a swap file with 128 MB swap at `/var/lib/swap/swapfile`, use the commands:

```
mkdir -p /var/lib/swap
dd if=/dev/zero of=/var/lib/swap/swapfile bs=1M count=128
```

- 2** Initialize this swap file with the command

```
mkswap /var/lib/swap/swapfile
```

- 3** Activate the swap with the command

```
swapon /var/lib/swap/swapfile
```

To disable this swap file, use the command

```
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
```

15.1.7 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 the *Storage Administration Guide*.

15.2 LVM Configuration

This section briefly describes the principles behind the Logical Volume Manager (LVM) and its multipurpose features. In Section 15.2.2, “LVM Configuration with YaST” (page 289), learn how to set up LVM with YaST.

WARNING

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.

15.2.1 The Logical Volume Manager

The LVM enables flexible distribution of hard disk space over several file systems. It was developed because sometimes the need to change the segmenting of hard disk space arises just after the initial partitioning has been done. Because it is difficult to modify partitions on a running system, LVM provides a virtual pool (volume group, VG for short) of memory space from which logical volumes (LVs) can be created as needed. The operating system accesses these LVs instead of the physical partitions. Volume groups can occupy more than one disk, so that several disks or parts of them may constitute one single VG. This way, LVM provides a kind of abstraction from the physical disk space that allows its segmentation to be changed in a much easier and safer way than with physical repartitioning. Background information regarding

physical partitioning can be found in Section 15.1.1, “Partition Types” (page 277) and Section 15.1, “Using the YaST Partitioner” (page 275).

Figure 15.3: *Physical Partitioning versus LVM*




Figure 15.3, “Physical Partitioning versus LVM” (page 288) compares physical partitioning (left) with LVM segmentation (right). On the left side, one single disk has been divided into three physical partitions (PART), each with a mount point (MP) assigned so that the operating system can gain access. On the right side, two disks have been divided into two and three physical partitions each. Two LVM volume groups (VG 1 and VG 2) have been defined. VG 1 contains two partitions from DISK 1 and one from DISK 2. VG 2 contains the remaining two partitions from DISK 2. In LVM, the physical disk partitions that are incorporated in a volume group are called physical volumes (PVs). Within the volume groups, four LVs (LV 1 through LV 4) have been defined. They can be used by the operating system via the associated mount points. The border between different LVs do not need to be aligned with any partition border. See the border between LV 1 and LV 2 in this example.

LVM features:

- Several hard disks or partitions can be combined in a large logical volume.
- Provided the configuration is suitable, an LV (such as `/usr`) can be enlarged if free space is exhausted.
- With LVM, it is possible to add hard disks or LVs in a running system. However, this requires hot-swappable hardware.
- It is possible to activate a "striping mode" that distributes the data stream of a LV over several PVs. If these PVs reside on different disks, the read and write performance is enhanced, as with RAID 0.

- The snapshot feature enables consistent backups (especially for servers) of the running system.

With these features, LVM is ready for heavily used home PCs or small servers. LVM is well-suited for the user with a growing data stock (as in the case of databases, music archives, or user directories). This would allow file systems that are larger than the physical hard disk. Another advantage of LVM is that up to 256 LVs can be added. However, working with LVM is different from working with conventional partitions. Instructions and further information about configuring LVM is available in the official LVM HOWTO at <http://tldp.org/HOWTO/LVM-HOWTO/>.

Starting from Kernel version 2.6, LVM version 2 is available, which is backward-compatible with the previous LVM and enables the continued management of old volume groups. When creating new volume groups, decide whether to use the new format or the backward-compatible version. LVM 2 does not require any kernel patches. It makes use of the device mapper integrated in kernel 2.6. This kernel only supports LVM version 2. Therefore, when talking about LVM, this section always refers to LVM version 2.

15.2.1.1 Thin Provisioning

Starting from Kernel version 3.4, LVM supports thin provisioning. A thin-provisioned volume has a virtual capacity and a real capacity. *Virtual* capacity is the volume storage capacity that is available to a host. *Real* capacity is the storage capacity that is allocated to a volume copy from a storage pool. In a fully allocated volume, the virtual capacity and real capacity are the same. In a thin-provisioned volume, however, the virtual capacity can be much larger than the real capacity. If a thin-provisioned volume does not have enough real capacity for a write operation, the volume is taken offline and an error is logged.

For more general information, see http://wikibon.org/wiki/v/Thin_provisioning.

15.2.2 LVM Configuration with YaST

The YaST LVM configuration can be reached from the YaST Expert Partitioner (see Section 15.1, “Using the YaST Partitioner” (page 275)) within the *Volume Management* item in the *System View* pane. The Expert Partitioner allows you to

edit and delete existing partitions and also create new ones that need to be used with LVM. The first task is to create PVs 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.

15.2.2.1 Creating Volume Groups

If no volume group exists on your system, you must add one (see Figure 15.4, “Creating a Volume Group” (page 291)). It is possible to create additional groups by clicking on *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 on *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 15.4: *Creating a Volume Group*

Add Volume Group

Volume Group Name
system

Physical Extent Size
4 MB

Available Physical Volumes:

Device	Size	Enc	Type
--------	------	-----	------

Selected Physical Volumes:

Device	Size	Enc	Type
/dev/sdb1	8.00 GB		Linux LVM

Total size: 0 B Resulting size: 8.00 GB

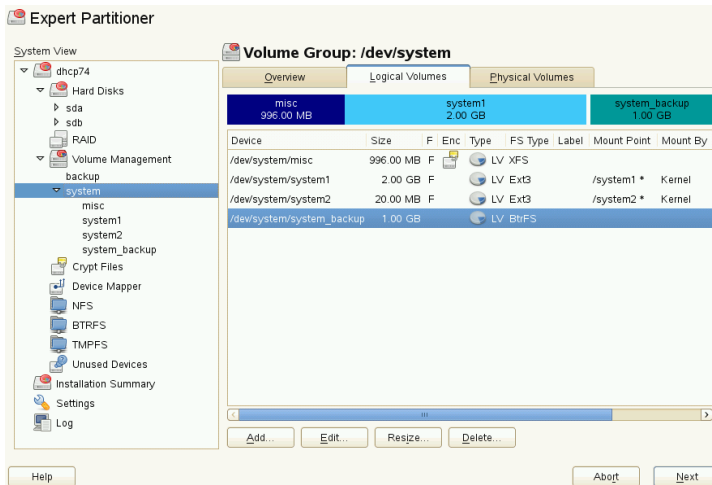
Help Abort Back Finish

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 or remove PVs to the selected volume group.

15.2.2.2 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 15.5: *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 self-explanatory 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.
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.

TIP

The big advantage of a thin provisioning is that the total sum of all thin volumes stored in a thin pool can exceed the size of the pool itself.

4. Choose the file system to use on the LV as well as 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, at this point, verify the correctness of your entries concerning striping. Any mistake made here is apparent only 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.

15.3 Soft RAID Configuration

The purpose of RAID (redundant array of independent disks) is to combine several hard disk partitions into one large *virtual* hard disk to optimize performance and/or data security. Most RAID controllers use the SCSI protocol because it can address a larger number of hard disks in a more effective way than the IDE protocol. It is also more suitable for the parallel command processing. There are some RAID controllers that support IDE or SATA hard disks. Soft RAID provides the advantages of RAID systems without the additional cost of hardware RAID controllers. However, this requires some CPU time and has memory requirements that make it unsuitable for high performance computers.

With SUSE® Linux Enterprise Server , you can combine several hard disks into one soft RAID system. RAID implies several strategies for combining several hard disks in a RAID system, each with different goals, advantages, and characteristics. These variations are commonly known as *RAID levels*.

Common RAID levels are:

RAID 0

This level improves the performance of your data access by spreading out blocks of each file across multiple disk drives. Actually, this is not really a RAID, because it does not provide data backup, but the name *RAID 0* for this type of system is commonly used. With RAID 0, two or more hard disks are pooled together. Performance is enhanced, but the RAID system is destroyed and your data lost if even one hard disk fails.

RAID 1

This level provides adequate security for your data, because the data is copied to another hard disk 1:1. This is known as *hard disk mirroring*. If one disk is destroyed, a copy of its contents is available on the other one. All disks but one could be damaged without endangering your data. However, if the damage is not detected, the damaged data can be mirrored to the undamaged disk. This could result in the same loss of data. The writing performance suffers in the copying process compared to using single disk access (10 to 20 % slower), but read access is significantly faster in comparison to any one of the normal physical hard disks. The reason is that the duplicate data can be parallel-scanned. Generally it can be said that Level 1 provides nearly twice the read transfer rate of single disks and almost the same write transfer rate as single disks.

RAID 5

RAID 5 is an optimized compromise between Level 0 and Level 1, in terms of performance and redundancy. The hard disk space equals the number of disks used minus one. The data is distributed over the hard disks as with RAID 0. *Parity blocks*, created on one of the partitions, exist for security reasons. They are linked to each other with XOR, enabling the contents to be reconstructed by the corresponding parity block in case of system failure. With RAID 5, no more than one hard disk can fail at the same time. If one hard disk fails, it must be replaced as soon as possible to avoid the risk of losing data.

RAID 6

To further increase the reliability of the RAID system, it is possible to use RAID 6. In this level, even if two disks fail, the array still can be reconstructed. With RAID 6, at least 4 hard disks are needed to run the array. Note that when running as software raid, this configuration needs a considerable amount of CPU time and memory.

RAID 10 (RAID 1+0)

This RAID implementation combines features of RAID 0 and RAID 1: the data is first mirrored to separate disk arrays, which are inserted into a new RAID 0; type array. In each RAID 1 sub-array, one disk can fail without any damage to the data. A minimum of four disks and an even number of disks is needed to run a RAID 10. This type of RAID is used for database application where a huge load is expected.

Other RAID Levels

Several other RAID levels have been developed (RAID 2, RAID 3, RAID 4, RAIDn, RAID 10, RAID 0+1, RAID 30, RAID 50, etc.), some of them being

proprietary implementations created by hardware vendors. These levels are not very common and therefore are not explained here.

15.3.1 Soft RAID Configuration with YaST

The YaST *RAID* configuration can be reached from the YaST Expert Partitioner, described in Section 15.1, “Using the YaST Partitioner” (page 275). 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.

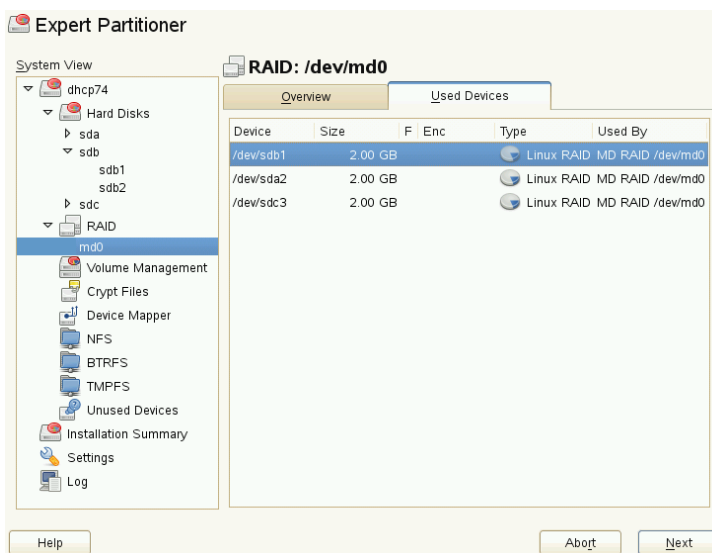
In the next dialog, choose between RAID levels 0, 1, 5, 6 and 10. Then, select all partitions with either the “Linux RAID” or “Linux native” type that should be used by the RAID system. No swap or DOS partitions are shown.

TIP

For RAID types where the order of added disks matters, you can mark individual disks with one of the letters A to E. Click the *Classify* button, select the disk and click one of the *Class X* buttons, where X is the letter

you want to assign to the disk. Assign all available RAID disks this way, and confirm with *OK*. You can easily sort the classified disks with the *Sorted* or *Interleaved* buttons, or add a sort pattern from a text file with *Pattern File*.

Figure 15.6: 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 as well as 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.

15.3.2 Troubleshooting

Check the file `/proc/mdstat` to find out whether a RAID partition has been damaged. In the event of a system failure, shut down your Linux system and replace the defective hard disk with a new one partitioned the same way. Then restart your system and enter the command `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.

15.3.3 For More Information

Configuration instructions and more details for soft RAID can be found in the HOWTOs at:

- `/usr/share/doc/packages/mdadm/Software-RAID.HOWTO.html`
- <http://raid.wiki.kernel.org>

Linux RAID mailing lists are available, such as <http://marc.info/?l=linux-raid>.

Subscription Management

Any machine running SUSE Linux Enterprise Server 11 or SUSE Linux Enterprise Desktop 11 can be configured to register against local Subscription Management Tool server to download software updates instead of communicating directly with the Novell Customer Center and the NU servers. To use an SMT server for client registration and as a local update source, you must configure the SMT server in your network first. The SMT server software is distributed as an add-on for SUSE Linux Enterprise Server and its configuration is described in the *Subscription Management Tool Guide*. There is no need to install any add-on on the clients to be configured for registering against an SMT server.

To register a client against an SMT server, you need to equip the client with the server's URL. As client and server communicate via the HTTPS protocol during registration, you also need to make sure the client trusts the server's certificate. In case your SMT server is set up to use the default server certificate, the CA certificate will be available on the SMT server via HTTP protocol at `http://FQDN/smt.crt`. In this case you do not have to concern yourself with the certificate: the registration process will automatically download the CA certificate from there, unless configured otherwise. You must enter a path to the server's CA certificate if the certificate was issued by an external certificate authority.

NOTE: Registering Against *.novell.com Subdomain

If you try to register against any *.novell.com subdomain, the certificate will not be downloaded during registration (for security reasons), and certificate handling will not be done. In such cases, use a different domain name or a plain IP address.

There are several ways to provide this information and to configure the client machine to use SMT. The first way is to provide the needed information via kernel parameters at boot time. The second way is to configure clients using an AutoYaST profile. There is also a script distributed with Subscription Management Tool, `clientSetup4SMT.sh`, which can be run on a client to make it register against a specified SMT server. These methods are described in the following sections:

16.1 Using Kernel Parameters to Access an SMT Server

Any client can be configured to use SMT by providing the following kernel parameters during machine boot: `regurl` and `regcert`. The first parameter is mandatory, the latter is optional.

`regurl`

URL of the SMT server. The URL needs to be in the following format: `https://FQDN/center/regsvc/` with *FQDN* being the fully qualified hostname of the SMT server. It must be identical to the FQDN of the server certificate used on the SMT server. Example:

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

`regcert`

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

URL

Remote location (http, https or ftp) from which the certificate can be downloaded. Example:

```
regcert=http://smt.example.com/smt.crt
```

Floppy

Specifies a location on a floppy. The floppy has to be inserted at boot time (you will not be prompted to insert it if it is missing). The value must start with the string `floppy`, followed by the path to the certificate. Example:

```
regcert=floppy/smt/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 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 either the certificate will be installed by an add-on product, or if you are using a certificate issued by an official certificate authority.

Example:

```
regcert=done
```

WARNING: Beware of Typing Errors

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

If a wrong value for `regcert` has been entered, you will be prompted for a local path to the certificate. In case `regcert` is not specified at all, it will default to `http://FQDN/smt.crt` with `FQDN` being the name of the SMT server.

WARNING: Change of SMT Server Certificate

If the SMT server gets a new certificate from a new and untrusted CA, the clients need to fetch the new CA certificate file. This is done automatically with the registration process but only if a URL was used at installation time to retrieve the certificate, or if the `regcert` parameter was omitted and thus, the default URL is used. If the certificate was loaded using any other method (such as floppy or local path), the CA certificate will not be updated.

16.2 Configuring Clients Using AutoYaST Profile

Clients can be configured to register with SMT server via AutoYaST profile. For general information about creating AutoYaST profiles and preparing automatic installation, refer to Chapter 21, *Automated Installation* (page 337). In this section, only SMT specific configuration is described.

To configure SMT specific data using AutoYaST, follow these steps:

- 1 As `root`, start YaST and select *Miscellaneous > Autoinstallation* to start the graphical AutoYaST front-end.

From a command line, you can start the graphical AutoYaST front-end with the `yast2 autoyast` command.

- 2 Open an existing profile using *File > Open*, create a profile based on the current system's configuration using *Tools > Create Reference Profile*, or just work with an empty profile.
- 3 Select *Support > Novell Customer Center Configuration*. An overview of the current configuration is shown.
- 4 Click *Edit*.
- 5 To register while installing automatically, select *Run Product Registration*. You can include information from your system with *Hardware Profile* and *Optional Information*.
- 6 Set the URL of the *SMT Server* and, optionally, the location of the *SMT Certificate*. The possible values are the same as for the kernel parameters `regurl` and `regcert` (see Section 16.1, “Using Kernel Parameters to Access an SMT Server” (page 300)). The only exception is that the `ask` value for `regcert` does not work in AutoYaST, because it requires user interaction. If using it, the registration process will be skipped.
- 7 Perform all other configuration needed for the systems to be deployed.
- 8 Select *File > Save As* and enter a filename for the profile, such as `autoinst.xml`.

16.3 Configuring Clients Using the `clientSetup4SMT.sh` Script

The `/usr/share/doc/packages/smt/clientSetup4SMT.sh` script is provided with SMT. This script allows to configure a client machine to use a SMT server or to reconfigure it to use a different SMT server.

To configure a client machine to use SMT with the `clientSetup4SMT.sh` script, follow these steps:

- 1 Copy the `/usr/share/doc/packages/smt/clientSetup4SMT.sh` script from your SMT server to the client machine.
- 2 As `root`, execute the script on the client machine. The script can be executed in two ways. In the first case, the script name is followed by the registration URL: `./clientSetup4SMT.sh registration_URL`, for example, `./clientSetup4SMT.sh https://smt.example.com/center/regsvc`. In the second case, the script name is followed by the `--host` option followed by hostname of the SMT server: `./clientSetup4SMT.sh --host server_hostname`, for example, `./clientSetup4SMT.sh --host smt.example.com`.
- 3 The script downloads the server's CA certificate. Accept it by pressing `y`.
- 4 The script performs all necessary modifications on the client. However, the registration itself is not performed by the script.
- 5 Perform a registration by executing `suse_register` or running `yast2 inst_suse_register` module on the client.

16.4 Registering Clients Against SMT Test Environment

To configure a client to register against the test environment instead the production environment, modify `/etc/suseRegister.conf` on the client machine by setting:

```
register = command=register&testenv=1
```

For more information about using SMT with a test environment, refer to the *Subscription Management Tool Guide*.

Part III. Imaging and Creating Products

KIWI

KIWI is a system for creating operating system images. An image is a directory with a file containing the operating system, its applications and configurations, the file system structure of the OS, possible additional metadata, and (depending on the image type) also disk geometry and partition table data. With KIWI you can create LiveCDs and LiveDVDs, USB sticks, virtual disk to play in full virtual systems like VMware, XEN images for paravirtualization in a hypervisor, and a PXE environment to boot from network.

17.1 Prerequisites for KIWI

To build images with KIWI, you need the following preconditions:

1. Free sufficient disk space for the operation.
2. KIWI is split into several packages, targeted to different image types. In any case, you need the base package `kiwi`. Depending on your target image, you need the following packages:

Image Type	Package Name
Installation Media	<code>kiwi-desc-oemboot</code>
Virtualization	<code>kiwi-desc-xenboot</code>

Image Type	Package Name
USB Sticks	kiwi-desc-usbboot
Network Client	kiwi-desc-netboot

3. Install the `kiwi-doc` package. You can find some example configurations to get an idea about structure and content.
4. Know the KIWI configuration file and its structure. It is based on a RELAX NG schema and documented in the `kiwi` package under `/usr/share/doc/packages/kiwi/kiwi.html`. You need this document, if you want to create the configuration file from scratch or when you want to insert elements or attributes.

17.2 Knowing KIWI's Build Process

The building process of KIWI is separated into three steps:

1. **Physical Extend (Preparation)** This stage prepares the content of your new file system. During this step the root directory is created, you determine which packages are installed on your image and which user configuration files are included.
2. **Logical Extend (Creation)** This stage requires a successful preparation step. The logical extend step creates the operating system image based on the first step.
3. **Deployment** The resulting image type can be deployed with different methods like installed on a hard disk or played by a virtualization system (VMware, Qemu, VirtualBox).

17.3 Image Description

KIWI needs an image description to build an image type. The image description is a directory which contains at least a file `config.xml`, or alternatively with the extension `*.kiwi`.

17.3.1 Contents of Image Description

The following table contains additional optional information. However, most of this information is mandatory for the later functionality of the operating system:

Table 17.1: *Additional Files and Directories For Image Description*

File/Directory	Description
config/	optional subdirectory. Contains Bash scripts which are executed after the installation of all the image packages.
config.sh	optional configuration script while creating the physical extend
config.xml	configuration file for each image description, explained in Section 17.3.2 (page 310)
config-cdroot.tgz	archive, only used for ISO images
config-cdroot.sh	manipulate extracted data from config-cdroot.tgz
config-yast-autoyast.xml	configuration file created by AutoYaST
config-yast-firstboot.xml	configuration file for controlling the YaST firstboot service
images.sh	optional configuration script while creating the preparation step
root/	contains other directories, special files, and scripts which are changed <i>after</i> the installation of all image packages

17.3.2 The config.xml File

All information about an image description is stored in the central configuration XML file `config.xml`. Each time KIWI is executed, `config.xml` is validated against an RELAX NG schema (see <http://www.relaxng.org> for more information about this schema language). Therefore, it is recommended to use an adequate XML editor with RELAX NG support or to use the documentation about the schema in the HTML file `/usr/share/doc/packages/kiwi/schema/kiwi.xsd.html`.

The configuration file consists of several parts:

- some description about the author, contact information, and a short explanation.
- preferences option needed for the logical extent stage.
- information about the users, their name, their home directories, and their passwords.
- links to repositories.
- a list of packages that are used for the defined image type.
- and other less important information which can be viewed in the above HTML file of the RELAX NG schema documentation.

A skeleton of the file is shown in the following example:

Example 17.1: *KIWI Configuration File*

```
<image schemeversion="2.0" name="..."> ❶
  <description type="system"> ❷
    <author>...</author>
    <contact>...</contact>
    <specification>...</specification>
  </description>
  <preferences> ❸
    <type primary="true" boot="..." flags="...">iso</type>
    <type boot="..." filesystem="ext3" format="vmkd">vmx</type>
    <type boot="..." filesystem="ext3">xen</type>
    <type boot="..." filesystem="squashfs" flags="unified">oem</type>
    <version>2.7.0</version>
    <size unit="M">780</size>
```

```

    <packagemanager>zypper</packagemanager>
    <rpm-check-signatures>False</rpm-check-signatures>
    <rpm-force>False</rpm-force>
    <locale>en_US.UTF-8</locale>
    <oem-swap>no</oem-swap>
    <oem-boot-title>USB</oem-boot-title>
</preferences>
<users group="users"> ❹
  <user name="root" pwd="" home="/root"/>
</users>
<repository type="rpm-md"> ❺
  <source path="/home/rpmdir"/>
</repository>
<packages type="image" patternPackageType="onlyRequired"> ❻
  <package name="yast2-live-installer"/>
  <package name="pam"/>
  <!-- List of packages reduced -->
</packages>

```

- ❶ The root element of every KIWI configuration file. Each file requires the version number. An optional `kiwirevision` attribute can be used to specify an SVN revision of KIWI.
- ❷ Contains mandatory descriptions with information about the creator of this image descriptions, its contact address and a short explanation.
- ❸ Contains mandatory preferences with information about the version of this image, the used package manager, the supported image types, and other settings.
- ❹ The optional `users` element contains a list of all users which are added to the image. The `user` element contains the name, the path to its home directory, password, and the shell.
- ❺ Contains a mandatory list of repositories used by the package manager.
- ❻ Contains a mandatory list of packages which are included into the image.

More details about the configuration file is shown in the HTML page above.

17.4 Creating Appliances with KIWI

This section describes how to create appliances with KIWI. An appliance is an especially-designed operating system for a specific task. For example, you can create an appliance with the focus on office programs.

17.4.1 Creating a Local Installation Source

All examples in `kiwi-doc` packages need a valid installation source to create an image. Usually the examples connect to a network resource. The higher the network bandwidth, the faster the image creation. If you do not have a fast network or you do not want to use it, create a local installation resource. Proceed as follows:

- 1 Collect your installation DVD.
- 2 Open a shell and become `root`.
- 3 Create the directory for your local installation directory. The examples use usually the path `/image/CDs/full-VERSION-ARCH`. Replace the placeholders `VERSION` and `ARCH` with the respective values.

- 4 Mount the medium. Replace the `DRIVE` placeholder with the respective device (usually `dvd`, `cdrom`, etc.):

```
mount -o loop /dev/DRIVE /mnt
```

- 5 Copy all the content of the medium into the installation directory:

```
cp -a /mnt/* /images/CDs/full-VERSION-ARCH
```

To use the local installation source, all you need to do is to enable it in the `repository` element:

```
<repository type="...">  
  <!-- Remove the comment markers in the next line -->  
  <!-- <source path="/image/CDs/full-VERSION-ARCH" -->  
  <source path="opensuse://openSUSE:11.0/standard"/>  
</repository>
```

17.4.2 Creating an Image

An image is a virtual disk image containing all partitions, boot loader information, and packages as it resides on a real disk. To create an ISO image, proceed as follows:

- 1 Install the packages `kiwi` and `kiwi-doc` and resolve any dependencies.
- 2 Open a shell and become `root`.

- 3** Copy the directory `/usr/share/doc/packages/kiwi/examples/suse-11.0/suse-oem-preload` to your current directory.
- 4** Open the file `config.xml` and locate the element `repository`. If you want to use a local installation source, refer to Section 17.4.1 (page 312) for more information.
- 5** Execute KIWI with the following command to prepare the first stage (“physical extend”):

```
kiwi --prepare suse-oem-preload --root oem
```

- 6** Build the ISO image:

```
kiwi --create oem --type iso --destdir /tmp/myoem
```

17.4.3 Creating Preload Image with NFS

To create an image with NFS functionality, proceed as follows:

- 1** Open a shell and become `root`.
- 2** Copy the directory `/usr/share/doc/packages/kiwi/examples/suse-11.1/suse-oem-preload` to your current directory.
- 3** Open the file `suse-oem-preload/config.xml` and locate the `packages` element with the attribute `type="image"`.
- 4** Insert the following line between `<packages type="image">` and `</packages>` and save the file:

```
<package name="nfs-client"/>
```
- 5** Rebuild the image as described in Step 5 (page 313).

17.5 For More Information

Find more information about KIWI in the following documents:

- <http://developer.berlios.de/projects/kiwi>—Homepage of KIWI

- Extensive description about the KIWI Image System: `/usr/share/doc/packages/kiwi/kiwi.pdf`

Creating Add-On Products With Add-on Creator

18

An Add-On is a special designed media, usually a CD or DVD, to extend your product. The Add-on Creator was developed to support our customers and partners and simplify third-party software distribution for all SUSE products.

18.1 Creating Images

To create an Add-On CD, proceed as follows:

- 1** Start YaST and open the *Add-On Creator* module. A window opens.
- 2** If you have not run this module before, click on *Create an Add-On from the Beginning* to start. In case you have already created an Add-On, the window shows a list of all created Add-Ons. Click *Add* to start.
- 3** Enter the product name and version of your Add-On and give some further options:
 - Choose the required product upon which it is based.
 - Select the path to additional Add-On packages. You need this, if you need further RPM packages which are not included in your base product (this step is optional).
 - Select the path with the required product packages (this step is optional).

- 4 Correct the product definition and enter a vendor name. Disable *Show Only Required Keywords* to display more keywords.
- 5 Change the package descriptions. Use *Add Language* to insert a new language and add translated descriptions (this step is optional).
- 6 Add new patterns. With patterns you can group your RPM packages. Use *New* to add a new pattern name and change the respective attributes in the list below (this step is optional).
- 7 Modify the output settings. Enter a path to your output directory and change the name of the ISO name (changing the name of the ISO is optional). Additionally, you can modify further features:
 - Use *Configure Workflow...* to enter files to customize your product workflow.
 - Use *Optional Files...* to add files to your Add-On product. The first part can be used to insert information about the Add-On in the `info.txt` file. Use the license files to display a window with *Agree* and *Disagree* buttons before the installation starts. More files can be added in the README section.

The second part can be used to store `COPYRIGHT` and `COPYING` files in various languages.
- 8 Sign your Add-On product with your GPG key. Signing your product with your GPG key provides evidence of the origin of your product. If you do not have a key, create one first and enter the respective passphrase twice.
- 9 Check your product in the overview and proceed with *Finish*.
- 10 Use the *Build* button to start the process. *Finish* closes the window.

18.2 Add-On Structure

If you create an Add-On product, the following overview contains the structure of the files and directories:

`ARCHIVES.gz`

Contains the gzipped contents of all RPM files. It is actually a listing of the `rpm` command with the options `-qil` for each RPM file.

`Changelog`

Contains all the changes of the RPM files.

`content`

Contains information about your Add-On product.

`content.asc`

Contains the signature file from GPG.

`content.key, gpg-pubkey-NUMBER.asc`

The public GPG key.

`INDEX.gz`

Contains a list of all RPM files and packed with `gzip`.

`ls-lR.gz`

Contains a list of all files and directories of your Add-On product medium.

`media.N/`

Contains files with basic information about the Add-On media set. The directory is numbered, so `media.1/` is for the first Add-On medium. Additional media have a consecutive number.

`suse/`

Contains sub directories with architecture-specific information. Exceptions are `noarch/` for architecture-independent packages, and `src/` for source packages. Proprietary software packages are stored under `nosrc/`.

18.3 For More Information

Find more information in the following documents:

- <http://old-en.opensuse.org/KIWI>—KIWI Project
- http://old-en.opensuse.org/Creating_YaST_Installation_Sources—YaST installation source creation
- http://old-en.opensuse.org/Standards/YaST2_repository_metadata—YaST metadata description

- [http://developer.novell.com/wiki/index.php/
Creating_Add-ons](http://developer.novell.com/wiki/index.php/Creating_Add-ons)—

Creating Images with YaST Product Creator

19

The YaST Product Creator is a unified graphical front-end for KIWI and Add-on Creator. It was developed to provide image creation functionality in one place. All tools integrated in the YaST Product Creator are also available as separate YaST modules or applications.

19.1 Prerequisites for Product Creator

Before you can create images with the YaST Product Creator, make sure you meet the following prerequisites:

1. Install the package `yast2-product-creator` from the SDK under http://www.novell.com/developer/sle_sdk.html. This package needs other packages. Make sure you fulfill all dependencies.
2. Free sufficient disk space for the operation.

19.2 Creating Images

The Product Creator uses KIWI to create an image of a product. In case you are interested in manually developing such images, refer to Chapter 17, *KIWI* (page 307).

To create an image, proceed as follows:

- 1 If you are starting the Product Creator for the first time, enter the configuration name and choose the method for adding packages to the ISO image.

If you have used the Product Creator already, select *Add* to create a new product definition and enter the configuration name and choose the method.

- 2 Select or deselect package sources. To select a source, select it from the table and click *Select*. With *Create New...* execute the Add-on Creator, see Chapter 18, *Creating Add-On Products With Add-on Creator* (page 315) for more information. To add a different source, add the source in the YaST *Installation Sources* module first then run the Product Creator again. After source selection, click *Next*.

NOTE: Unsupported Target Architectures

Do not change the target architecture. KIWI does not presently support the building of different architectures.

- 3 Enter the path in which to create the skeleton directory. Choose whether to *Generate ISO Image File* or *Create Directory Tree Only*. Use the other options to insert metadata. Click *Next*.
- 4 Edit the content of the `isolinux.cfg` file, if it is a part of the configuration. In most cases you can leave it as it is. If the file is not part of the configuration, add it now with *Load File*. Click *Next*.
- 5 Select your software. All package dependencies are solved automatically after *Apply* is clicked.
- 6 Sign your product with *Digitally Sign the Product on the Medium*, if needed. Provide a key for your product configuration. Signing your product with your GPG key provides evidence of the origin of your product. After key configuration, click *Next*.
- 7 Review the summary. To change any option, use *Back*. To confirm your new product configuration, click *Finish*.

Your product definition is now completed. The Product Creator allows you to choose from the following actions:

- **Create Product** Creates an ISO image of the selected product. If there is something missing, the process will be aborted. Correct the error and repeat the configuration.
- **Create Image with KIWI...** Use the pull-down menu to choose from different target formats, such as Live media or Xen images.

19.3 For More Information

Find more information about creating system images and related topics in the following documents:

- Chapter 17, *KIWI* (page 307)
- <http://old-en.opensuse.org/KIWI>—The KIWI project
- KIWI documentation: `/usr/share/doc/packages/kiwi/kiwi.pdf`

Deploying Customized Preinstallations

20

Rolling out customized preinstallations of SUSE Linux Enterprise Server to a large number of 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; for more information, see Chapter 21, *Automated Installation* (page 337)).

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” (page 324).
- 2** Customize the firstboot workflow. For more information, refer to Section 20.2, “Customizing the Firstboot Installation” (page 324).
- 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” (page 333).
- 4** Have the end user personalize the instance of SUSE Linux Enterprise Server. For more information, refer to Section 20.4, “Personalizing the Installation” (page 333).

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 wait for the installed machine to boot. Also install the `yast2-firstboot` package.
- 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” (page 324). Otherwise proceed directly to Step 5 (page 324).
- 5 Enable firstboot as `root`:

Create an empty file `/var/lib/YaST2/reconfig_system` to trigger firstboot's execution. This file will be deleted once 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” (page 333).

20.2 Customizing the Firstboot Installation

Customizing the firstboot installation workflow may involve several different components. Customizing them is optional. 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” (page 325).
- Customizing licenses and license actions, as described in Section 20.2.2, “Customizing the License Action” (page 326).

- Customizing the release notes to display, as described in Section 20.2.3, “Customizing the Release Notes” (page 327).
- Customizing the order and number of components involved in the installation, as described in Section 20.2.4, “Customizing the Workflow” (page 327).
- Configuring additional optional scripts, as described in Section 20.2.5, “Configuring Additional Scripts” (page 332).

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” (page 332).

If you want to customize more than just the workflow components, refer the `control.xml` documentation at http://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:

- 2a** 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/"
```

- 2b** If your welcome message has filenames other than `welcome.txt` and `welcome_locale.txt` (where *locale* matches the ISO 639 language codes such as “cs” or “de”), specify the filename 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 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 if you have changed the instance of SUSE Linux Enterprise Server you are deploying with firstboot, you probably 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 filename 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

By default, a standard firstboot workflow includes the following components:

- Language Selection
- Welcome
- License Agreement
- Host Name
- Network
- Time and Date
- Desktop
- root Password

- User Authentication Method
- User Management
- Hardware Configuration
- Finish Setup

This standard layout of a firstboot installation workflow is not mandatory. You can enable or disable certain components or integrate your own modules into the workflow. To modify the firstboot workflow, manually edit 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.

For an overview about proposals, see Example 20.1, “Configuring the Proposal Screens” (page 328). 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 with 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`.

- ⑤ 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`.
- ⑥ The label to be displayed on the proposal.
- ⑦ The container for all modules that are part of the proposal screen.
- ⑧ 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” (page 328)). 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` and have the `.ycp` file suffix.

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 to 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 when 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.ycp` 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 possible dependencies to 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>
```

- 4a Enter the label to be displayed on your module in the `label` element.
- 4b Make sure that `enabled` is set to `true` to have your module included in the workflow.
- 4c Enter the filename of your module in the `name` element. Omit the full path and the `.ycp` suffix.

5 Apply your settings and close the configuration file.

TIP: Finding Connected Network Interface For Auto-Configuration

If the target hardware may feature more than one network interface add the `network-autoconfig` package to the application image. `network-autoconfig` makes sure that during firstboot all available ethernet interfaces are cycled until one is successfully configured with 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” (page 327). This is different from the localization of customized YaST messages, which is already described in Section 20.2.1, “Customizing YaST Messages” (page 325).

If you have changed `/etc/YaST2/firstboot.xml` and introduced string changes, generate a new translation template file (`.pot` file) and use the `gettext` tool chain 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 Change the `textdomain` setting from:

```
<textdomain>firstboot</textdomain>
```

to, for example,

```
<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 just changed an existing module, make sure to change also its text domain statement to avoid undesired side effects.

TIP: For More Information

For more information about YaST development, refer to http://en.opensuse.org/openSUSE:YaST_development. Detailed information about YaST firstboot can be found at <http://doc.opensuse.org/projects/YaST/SLES11/tdg/bk09ch01s02.html>.

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 Chapter 17, *KIWI* (page 307).

20.4 Personalizing the Installation

As soon as the cloned disk image is booted, firstboot starts and the installation proceeds exactly as laid out in Section 20.2.4, “Customizing the Workflow” (page 327). 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. Once this process is finished, a firstboot installed system behaves as any other instance of SUSE Linux Enterprise Server.

Part IV. Automated Installations

Automated Installation

AutoYaST allows you to install SUSE® Linux Enterprise on a large number of machines in parallel. The AutoYaST technology offers great flexibility to adjust deployments to heterogeneous hardware. This chapter tells you how to prepare a simple automated installation and lay out an advanced scenario involving different hardware types and installation purposes.

21.1 Simple Mass Installation

IMPORTANT: Identical Hardware

This scenario assumes you are rolling out SUSE Linux Enterprise to a set of machines with exactly the same hardware configuration.

To prepare for an AutoYaST mass installation, proceed as follows:

- 1 Create an AutoYaST profile that contains the installation details needed for your deployment as described in Section 21.1.1, “Creating an AutoYaST Profile” (page 338).
- 2 Determine the source of the AutoYaST profile and the parameter to pass to the installation routines as described in Section 21.1.2, “Distributing the Profile and Determining the autoyast Parameter” (page 340).
- 3 Determine the source of the SUSE Linux Enterprise installation data as described in Section 21.1.3, “Providing the Installation Data” (page 342).

- 4 Determine and set up the boot scenario for autoinstallation as described in Section 21.1.4, “Setting Up the Boot Scenario” (page 343).
- 5 Pass the command line to the installation routines by adding the parameters manually or by creating an `info` file as described in Section 21.1.5, “Creating the `info` File” (page 345).
- 6 Start the autoinstallation process as described in Section 21.1.6, “Initiating and Monitoring the Autoinstallation” (page 348).

21.1.1 Creating an AutoYaST Profile

An AutoYaST profile tells AutoYaST what to install and how to configure the installed system to get a completely ready-to-use system in the end. It can be created in several different ways:

- Clone a fresh installation from a reference machine to a set of identical machines
- Use the AutoYaST GUI to create and modify a profile to meet your requirements
- Use an XML editor and create a profile from scratch

To clone a fresh reference installation, proceed as follows:


- 1 Perform a normal installation.
- 2 After you complete the hardware configuration and read the release notes, check *Clone This System for AutoYaST*, if it is not yet checked by default. This creates a ready-to-use profile as `/root/autoyast.xml` that can be used to create clones of this particular installation.

To use the AutoYaST GUI to create a profile from an existing system configuration and modify it to your needs, proceed as follows:

- 1 As `root`, start YaST.
- 2 Select *Miscellaneous > Autoinstallation* to start the graphical AutoYaST front-end.
- 3 Select *Tools > Create Reference Profile* to prepare AutoYaST to mirror the current system configuration into an AutoYaST profile.

- 4 As well as the default resources (like boot loader, partitioning, and software selection), you can add various other aspects of your system to the profile by checking the items in the list in *Create a Reference Control File*.
- 5 Click *Create* to have YaST gather all the system information and write it to a new profile.
- 6 To proceed, choose one of the following:
 - If the profile is complete and matches your requirements, select *File > Save as* and enter a filename for the profile, such as `autoyast.xml`.
 - Modify the reference profile by selecting the appropriate configuration aspects (such as “Hardware/Printer”) from the tree view to the left and clicking *Configure*. The respective YaST module starts but your settings are written to the AutoYaST profile instead of applied to your system. When done, select *File > Save as* and enter a suitable name for the profile.
- 7 Leave the AutoYaST module with *File > Exit*.

Figure 21.1: *Editing an AutoYaST Profile with the AutoYaST Front-End*



21.1.2 Distributing the Profile and Determining the autoyast Parameter

The AutoYaST profile can be distributed in several different ways. Depending on the protocol used to distribute the profile data, different AutoYaST parameters are used to make the profile location known to the installation routines on the client. The location of the profile is passed to the installation routines by means of the boot prompt or an `info` file that is loaded upon boot. The following options are available:

Profile Location	Parameter	Description
File	<code>autoyast=file:// <i>path</i></code>	Makes the installation routines look for the control file in the specified path (relative to source root directory— <code>file:///autoyast.xml</code> if in the top directory of a CD-ROM).
Device	<code>autoyast=device:// <i>path</i></code>	Makes the installation routines look for the control file on a storage device. Only the device name is needed— <code>/dev/sda1</code> is wrong, use <code>sda1</code> instead.
Floppy	<code>autoyast=floppy:// <i>path</i></code>	Makes the installation routines look for the control file on a floppy in the floppy drive. This option is especially useful if you want to boot from CD-ROM.

Profile Location	Parameter	Description
NFS	<code>autoyast=nfs:// server/path</code>	Has the installation routines retrieve the control file from an NFS server.
HTTP	<code>autoyast=http:// server/path</code>	Has the installation routines retrieve the control file from an HTTP server.
HTTPS	<code>autoyast=https:// server/path</code>	Has the installation routines retrieve the control file from an HTTPS server.
TFTP	<code>autoyast=tftp:// server/path</code>	Has the installation routines retrieve the control file from a TFTP server.
FTP	<code>autoyast=ftp:// server/path</code>	Has the installation routines retrieve the control file from an FTP server.

Replace the *server* and *path* placeholders with values matching your actual setup.

AutoYaST includes a feature that allows the binding of certain profiles to the client's MAC address. Without having to alter the `autoyast=` parameter, you can have the same setup install several different instances using different profiles.

To use this, proceed as follows:

- 1 Create separate profiles with the MAC address of the client as the filename and put them on the HTTP server that holds your AutoYaST profiles.

- 2 Omit the exact path including the filename when creating the `autoyast=` parameter, for example:

```
autoyast=tftp://192.168.1.115/
```

- 3 Start the autoinstallation.

YaST tries to determine the location of the profile in the following way:

1. YaST searches for the profile using its own IP address in uppercase hexadecimal, for example, `192.0.2.91` is `C000025B`.
2. If this file is not found, YaST removes one hex digit and tries again. This action is repeated eight times until the file with the correct name is found.
3. If that still fails, it tries locating a file with the MAC address of the clients as the filename. The MAC address of the example client is `0080C8F6484C`.
4. If the MAC address-named file cannot be found, YaST searches for a file named `default` (in lowercase). An example sequence of addresses where YaST searches for the AutoYaST profile looks as follows:

```
C000025B
C000025
C00002
C0000
C000
C000
C00
C00
C0
C
0080C8F6484C
default
```

21.1.3 Providing the Installation Data

The installation data can be provided by means of the product CDs or DVDs or using a network installation source. If the product CDs are used as the installation source, physical access to the client to be installed is needed, because the boot process needs to be initiated manually and the CDs need to be changed.

To provide the installation sources over the network, set up a network installation server (HTTP, NFS, FTP) as described in Section 14.2.1, “Setting Up an Installation

Server Using YaST” (page 246). Use an `info` file to pass the server's location to the installation routines.

21.1.4 Setting Up the Boot Scenario

The client can be booted in several different ways:

Network Boot

As with a normal remote installation, autoinstallation can be initiated with Wake on LAN and PXE, the boot image and control file can be pulled in via TFTP, and the installation sources from any network installation server.

Bootable CD-ROM

You can use the original SUSE Linux Enterprise media to boot the system for autoinstallation and pull in the control file from a network location or a floppy. Alternately, create your own custom CD-ROM holding both the installation sources and the AutoYaST profile.

The following sections provide a basic outline of the procedures for network boot or boot from CD-ROM.

21.1.4.1 Preparing for Network Boot

Network booting with Wake on LAN, PXE, and TFTP is discussed in Section 14.1.3, “Remote Installation via VNC—PXE Boot and Wake on LAN” (page 241). To make the setup introduced there work for autoinstallation, modify the featured PXE Linux configuration file (`/srv/tftp/pxelinux.cfg/default`) to contain the `autoyast` parameter pointing to the location of the AutoYaST profile. An example entry for a standard installation looks like this:

```
default linux

# default label linux
    kernel linux
    append initrd=initrd install=http://192.168.1.115/install/suse-
enterprise/
```

The same example for autoinstallation looks like this:

```
default linux

# default label linux
```

```
kernel linux
append initrd=initrd install=http://192.168.1.115/install/suse-
enterprise/ \
    autoyast=nfs://192.168.1.110/profiles/autoyast.xml
```

Replace the example IP addresses and paths with the data used in your setup.

21.1.4.2 Preparing to Boot from CD-ROM

There are several ways in which booting from CD-ROM can come into play in AutoYaST installations. Choose from the following scenarios:

Boot from SUSE Linux Enterprise Media, Get the Profile over the Network

Use this approach if a totally network-based scenario is not possible (for example, if your hardware does not support PXE) and you have physical access to system to install during most of the process.

You need:

- The SUSE Linux Enterprise media
- A network server providing the profile data (see Section 21.1.2, “Distributing the Profile and Determining the autoyast Parameter” (page 340) for details)
- A floppy containing the `info` file that tells the installation routines where to find the profile

or

Access to the boot prompt of the system to install where you manually enter the `autoyast= parameter`

Boot and Install from SUSE Linux Enterprise Media, Get the Profile from a Floppy

Use this approach if an entirely network-based installation scenario would not work. It requires physical access to the system to be installed for turning on the target machine, or, in the second case, to enter the profile's location at the boot prompt. In both cases, you may also need to change media depending on the scope of installation.

You need:

- The SUSE Linux Enterprise media

- A floppy holding both the profile and the `info` file

or

Access to the boot prompt of the target to enter the `autoyast=` parameter

Boot and Install from Custom Media, Get the Profile from the Media

If you just need to install a limited number of software packages and the number of targets is relatively low, creating your own custom CD holding both the installation data and the profile itself might prove a good idea, especially if no network is available in your setup.

21.1.5 Creating the info File

The installation routines at the target need to be made aware of all the different components of the AutoYaST framework. This is done by creating a command line containing all the parameters needed to locate the AutoYaST components, installation sources, and the parameters needed to control the installation process.

Do this by manually passing these parameters at the boot prompt of the installation or by providing a file called `info` that is read by the installation routines (`linuxrc`). The former requires physical access to any client to install, which makes this approach unsuitable for large deployments. The latter enables you to provide the `info` file on some media that is prepared and inserted into the clients' drives prior to the autoinstallation. Alternatively, use PXE boot and include the `linuxrc` parameters in the `pxelinux.cfg/default` file as shown in Section 21.1.4.1, “Preparing for Network Boot” (page 343).

The following parameters are commonly used for `linuxrc`. For more information, refer to the AutoYaST package documentation under `/usr/share/doc/packages/autoyast`.

IMPORTANT: Separating Parameters and Values

When passing parameters to `linuxrc` at the boot prompt, use `=` to separate parameter and value. When using an `info` file, separate parameter and value with `:.`

Keyword	Value
<code>netdevice</code>	The network device to use for network setup (for BOOTP/DHCP requests). Only needed if several network devices are available.
<code>hostip</code>	When empty, the client sends a BOOTP request. Otherwise the client is configured using the specified data.
<code>netmask</code>	Netmask for the selected network.
<code>gateway</code>	Default gateway.
<code>nameserver</code>	Name server.
<code>autoyast</code>	Location of the the control file to use for the automatic installation, such as <code>autoyast=nfs://192.168.1.110/profiles/</code> .
<code>install</code>	Location of the installation source, such as <code>install=nfs://192.168.1.110/CDs/</code> .
<code>vnc</code>	If set to 1, enables VNC remote controlled installation.
<code>vncpassword</code>	The password for VNC.
<code>usessh</code>	If set to 1, enables SSH remote controlled installation.
<code>netsetup</code>	If set to 1, sets up the network. Normally this is done automatically, but you need to set <code>netsetup=1</code> in case the installation repository is

Keyword	Value
	provided locally (e.g. via DVD or local iso image) and the <code>info</code> file is loaded from the network.

If your autoinstallation scenario involves client configuration via DHCP and a network installation source, and you want to monitor the installation process using VNC, your `info` would look like this:

```
autoyast:profile_source install:install_source vnc:1
vncpassword:some_password
```

If you prefer a static network setup at installation time, your `info` file would look like the following:

```
autoyast:profile_source \
install:install_source \
hostip:some_ip \
netmask:some_netmask \
gateway:some_gateway
```

The `\` indicates that the line breaks have only been added for the sake of readability. All options must be entered as one continuous string.

The `info` data can be made available to `linuxrc` in various different ways:

- As a file on a floppy or CD Rom that is in the client's drive at installation time. Add the `info` parameter similar to `info=floppy:/info` or `info=cd:/info`.
- As a file in the root directory of the initial RAM disk used for booting the system provided either from custom installation media or via PXE boot.
- As part of the AutoYaST profile. In this case, the AutoYaST file needs to be called `info` to enable `linuxrc` to parse it. An example for this approach is given below.
- By means of an URL that points to the location of the `info` file. The syntax for this looks like `info=http://www.example.com/info`.

`linuxrc` looks for a string (`start_linuxrc_conf`) in the profile that represents the beginning of the file. If it is found, it parses the content starting from that string and finishes when the string `end_linuxrc_conf` is found. The options are stored in the profile as follows:

```

.....
    <install>
.....
    <init>
        <info_file>
<![CDATA[
#
# Don't remove the following line:
# start_linuxrc_conf
#
install: nfs:server/path
vnc: 1
vncpassword: test
autoyast: file:///info

# end_linuxrc_conf
# Do not remove the above comment
#
]]>

        </info_file>
    </init>
.....
    </install>
.....

```

linuxrc loads the profile containing the boot parameters instead of the traditional `info` file. The `install:` parameter points to the location of the installation sources. `vnc` and `vncpassword` indicate the use of VNC for installation monitoring. The `autoyast` parameter tells linuxrc to treat `info` as an AutoYaST profile.

21.1.6 Initiating and Monitoring the Autoinstallation

After you have provided all the infrastructure mentioned above (profile, installation source, and `info` file), you can go ahead and start the autoinstallation. Depending on the scenario chosen for booting and monitoring the process, physical interaction with the client may be needed:

- If the client system boots from any kind of physical media, either product media or custom CDs, you need to insert these into the client's drives.
- If the client is not switched on via Wake on LAN, you need to at least switch on the client machine.

- If you have not opted for remote controlled autoinstallation, the graphical feedback from AutoYaST is sent to the client's attached monitor or, if you use a headless client, to a serial console.

To enable remote controlled autoinstallation, use the VNC or SSH parameters described in Section 21.1.5, “Creating the `info` File” (page 345) and connect to the client from another machine as described in Section 14.5, “Monitoring the Installation Process” (page 270).

21.2 Rule-Based Autoinstallation

The following sections introduce the basic concept of rule-based installation using AutoYaST and provide an example scenario that enables you to create your own custom autoinstallation setup.

21.2.1 Understanding Rule-Based Autoinstallation

Rule-based AutoYaST installation allows you to cope with heterogeneous hardware environments:

- Does your site contain hardware of different vendors?
- Are the machines on your site of different hardware configuration (for example, using different devices or using different memory and disk sizes)?
- Do you intend to install across different domains and need to distinguish between them?

What rule-based autoinstallation does is, basically, generate a custom profile to match a heterogeneous scenario by merging several profiles into one. Each rule describes one particular distinctive feature of your setup (such as disk size) and tells AutoYaST which profile to use when the rule matches. Several rules describing different features of your setup are combined in an AutoYaST `rules.xml` file. The rule stack is then processed and AutoYaST generates the final profile by merging the different profiles matching the AutoYaST rules into one. To

illustrate this procedure, refer to Section 21.2.2, “Example Scenario for Rule-Based Autoinstallation” (page 351).

Rule-based AutoYaST offers you great flexibility in planning and executing your SUSE Linux Enterprise deployment. You can:

- Create rules for matching any of the predefined system attributes in AutoYaST
- Combine multiple system attributes (such as disk size and kernel architecture) into one rule by using logical operators
- Create custom rules by running shell scripts and passing their output to the AutoYaST framework. The number of custom rules is limited to five.

NOTE

For more information about rule creation and usage with AutoYaST, refer to the package's documentation under `/usr/share/doc/packages/autoyast2/html/index.html`, Chapter *Rules and Classes*.

To prepare for a rule-based AutoYaST mass installation, proceed as follows:

- 1 Create several AutoYaST profiles that contain the installation details needed for your heterogeneous setup as described in Section 21.1.1, “Creating an AutoYaST Profile” (page 338).
- 2 Define rules to match the system attributes of your hardware setup as shown in Section 21.2.2, “Example Scenario for Rule-Based Autoinstallation” (page 351).
- 3 Determine the source of the AutoYaST profile and the parameter to pass to the installation routines as described in Section 21.1.2, “Distributing the Profile and Determining the autoyast Parameter” (page 340).
- 4 Determine the source of the SUSE Linux Enterprise installation data as described in Section 21.1.3, “Providing the Installation Data” (page 342).
- 5 Pass the command line to the installation routines by adding the parameters manually or by creating an `info` file as described in Section 21.1.5, “Creating the `info` File” (page 345).

- 6 Determine and set up the boot scenario for autoinstallation as described in Section 21.1.4, “Setting Up the Boot Scenario” (page 343).
- 7 Start the autoinstallation process as described in Section 21.1.6, “Initiating and Monitoring the Autoinstallation” (page 348).

21.2.2 Example Scenario for Rule-Based Autoinstallation

To get a basic understanding of how rules are created, think of the following example, depicted in Figure 21.2, “AutoYaST Rules” (page 352). One run of AutoYaST installs the following setup:

A Print Server

This machine just needs a minimal installation without a desktop environment and a limited set of software packages.


Workstations in the Engineering Department

These machines need a desktop environment and a broad set of development software.

Laptops in the Sales Department

These machines need a desktop environment and a limited set of specialized applications, such as office and calendaring software.

Figure 21.2: *AutoYaST Rules*



In a first step, use one of the methods outlined in Section 21.1.1, “Creating an AutoYaST Profile” (page 338) to create profiles for each use case. In this example, you would create `print.xml`, `engineering.xml`, and `sales.xml`.

In the second step, create rules to distinguish the three hardware types from one another and to tell AutoYaST which profile to use. Use an algorithm similar to the following to set up the rules:

1. Does the machine have an IP of `192.168.2.253`? Then make it the print server.
2. Does the machine have PCMCIA hardware and feature an Intel chipset? Then consider it an Intel laptop and install the sales department software selection.
3. If none of the above is true, consider the machine a developer workstation and install accordingly.

Roughly sketched, this translates into a `rules.xml` file with the following content:

```
<?xml version="1.0"?>
```

```

<!DOCTYPE autoinstall SYSTEM "/usr/share/autoinstall/dtd/rules.dtd">
<autoinstall xmlns="http://www.suse.com/1.0/yast2ns" xmlns:config="http://
www.suse.com/1.0/configns">
  <rules config:type="list">
    <rule>
      <hostaddress>
        <match>192.168.2.253</match>
        <match_type>exact</match_type>
      </hostaddress>
      <result>
        <profile>print.xml</profile>
        <continue config:type="boolean">>false</continue>
      </result>
    </rule>
    <rule>
      <haspcmcia>
        <match>1</match>
        <match_type>exact</match_type>
      </haspcmcia>
      <custom1>
        <script>
if grep -i intel /proc/cpuinfo > /dev/null; then
echo -n "intel"
else
echo -n "non_intel"
fi;
        </script>
        <match>*</match>
        <match_type>exact</match_type>
      </custom1>
      <result>
        <profile>sales.xml</profile>
        <continue config:type="boolean">>false</continue>
      </result>
      <operator>and</operator>
    </rule>
    <rule>
      <haspcmcia>
        <match>0</match>
        <match_type>exact</match_type>
      </haspcmcia>
      <result>
        <profile>engineering.xml</profile>
        <continue config:type="boolean">>false</continue>
      </result>
    </rule>
  </rules>
</autoinstall>

```

When distributing the rules file, make sure that the rules directory resides under the profiles directory, specified in the `autoyast=protocol:serverip/profiles/URL`. AutoYaST looks for a rules subdirectory containing a file

named `rules.xml` first then loads and merges the profiles specified in the rules file.

The rest of the autoinstallation procedure is carried out as usual.

21.3 For More Information

For in-depth information about the AutoYaST technology, refer to *AutoYaST* ([↑AutoYaST](#)) or the documentation installed along with the software (`/usr/share/doc/packages/autoyast2`).

Automated Upgrade from SUSE Linux Enterprise 11 SP2 to 11 SP3

22

The following procedure is a way to do a mass upgrade unattended from SUSE Linux Enterprise 11 SP2 to SUSE Linux Enterprise 11 SP3. Several preparation steps are needed to create a suitable AutoYaST profile. AutoYaST finally will execute the upgrade process.

22.1 Preparing the AutoYaST Profile

The AutoYaST profile for the automated upgrade uses the same file format as the AutoYaST installation. For more information about AutoYaST, see Chapter 21, *Automated Installation* (page 337) and *AutoYaST* (↑*AutoYaST*).

However, for obvious reasons, there are some parts of the system (e.g., partitioning) that do not make sense to be configured during the upgrade. On the other hand, it is useful to set upgrade-specific options by means of the AutoYaST profile.

22.1.1 Upgrade

The upgrade options define the behavior of the dependency solver during upgrade:

```
<upgrade>
  <only_installed_packages
    config:type="boolean">false</only_installed_packages>
  <stop_on_solver_conflict
    config:type="boolean">true</stop_on_solver_conflict>
```

```
</upgrade>
```

`only_installed_packages`

Set to `true` for package-based upgrades (recommended for upgrading to the next service pack of the same product) or `false` for pattern-based upgrades (recommended for an upgrade between versions of a product, e.g. from SLES10 to SLES11).

`stop_on_solver_conflict`

Defines whether to show the proposal in case of failure to resolve package dependencies interactively (recommended to be set to `true`, but this could result in an interactive process, during which the user must to resolve the conflicts manually).

22.1.2 Software Selection

The software selection options define, which components to select or deselect in addition to the results of the resolver:

```
<software>
  <packages config:type="list">
    <package>autoyast2-installation</package>
    <package>apparmor-profile-editor</package>
  </packages>
  <patterns config:type="list">
    <pattern>base</pattern>
  </patterns>
  <remove-packages config:type="list"/>
  <remove-patterns config:type="list"/>
</software>
```

It is especially important to set packages or patterns for being selected or deselected in order to resolve package conflicts and thus to avoid the need for interactive intervention. Once the upgrade is done, the newly created `autoupg_updated.xml` file contains these packages and patterns plus those that were selected or deselected for any other reason.

22.1.3 Backup Before Upgrade

The backup before upgrade options match these features in the upgrade proposal.

```
<backup>
  <sysconfig config:type="boolean">true</sysconfig>
```

```
<modified config:type="boolean">true</modified>
<remove_old config:type="boolean">false</remove_old>
</backup>
```

`sysconfig`

defines whether to backup `sysconfig` before upgrading.

`modified`

defines whether to backup the modified configuration files before upgrading.

`remove_old`

defines whether to remove old backups from previous upgrades.

22.2 Running the Automated Upgrade

To start the automated upgrade, boot the installation media and pass the AutoYaST profile to it. There are two ways to pass the profile to the system:

- Pass the profile to the kernel command line the same way as for the AutoYaST installation (use the `autoupgrade=1 autoyast=http://host/path/profile.xml` parameter. For System z, this is the only possibility.
- Pass the `autoupgrade=1` parameter to the kernel command line. Before you start the upgrade, copy the profile to `/root/autoupg.xml`. Then there is no need for any additional kernel parameter.

The latter approach allows you to have a single installation kernel command line for even different machines—just copy the appropriate profile into its file system.

As long as you have only one SUSE Linux Enterprise system installed on your machine, there are no package conflicts and you did not set the profile to stop on the upgrade proposal, the complete process will be non-interactive. In case you enter the upgrade proposal, you can modify its settings for the upgrade.

After the upgrade finishes, YaST writes the `/root/autoupg-updated.xml` file, which contains the profile with applied software selection changes done in the proposal. This is especially useful in case of mass upgrades of machines with the same package selection. This way, conflict resolutions from one machine can easily

be applied on other machines, which consequently will get these conflicts resolved automatically and the upgrade itself will be non-interactive.

If there are more SUSE Linux Enterprise systems installed on the machine, you will always be asked, which one to upgrade—there is no way to select it in advance.

22.3 GRUB Menu Section for Booting into the Upgrade

An alternate way to boot the system is to create an additional section in the GRUB menu (and similar for other boot loaders and other architectures), which starts the installation. The following example assumes that there is a separate `/boot` partition, which is referred in GRUB as `(hd0, 0)`:

```
title Upgrade
    root (hd0,0)
    kernel /upgrade/linux
    install=inst_source_url autoupgrade=1
    autoyast=autoyast_profile_url vga=0x314
    initrd /upgrade/initrd
```

The above example assumes that the installation kernel and the installation `initrd` are located in the `/boot/upgrade` directory.

On System `z`, you must add the parameters to the `PARM` file—proceed the same way as you do when performing an AutoYaST-driven installation.

22.4 Second Stage of the Upgrade

The automated upgrade by default does not perform configuration changes during the second stage of the upgrade. The only exception is network configuration, which needs to be set to be preserved in the AutoYaST upgrade profile.

If configuration adjustment of some system areas is needed after the upgrade (e.g., configuring a new service), add the relevant sections to the AutoYaST profile for the upgrade and the configuration of the selected system areas will be saved during the upgrade.

WARNING: AutoYaST Supplied Configuration Replaces Existing Configuration

Be warned that existing configuration of that system area will be replaced and thus destroyed by the AutoYaST configuration.

Normally, the only configuration adjustment, which should be present in the AutoYaST profile, is the registration of the system with Subscription Management Tool (SMT) or Novell Customer Center (NCC). If this is missing, the system will not get the update repository and updates will not be possible—unless configured later again.

22.5 Limitations and Hints

22.5.1 NetworkManager and Registration

In case of using NetworkManager for managing network devices and network connections, network connection is not available during the second stage of the upgrade. This prevents the system from performing the registration.

22.5.2 Cleaning Up Upgrade Setting

If you do any changes in your system in order to trigger the upgrade process (e.g., adding a new section to the boot loader menu), you probably want to remove it after the upgrade is done.

You can do it automatically with a post-installation script. Find an examples in Section “Custom User Scripts” (Chapter 4, *Configuration and Installation Options*, ↑*AutoYaST*). A sample script cleaning up GRUB's `menu.lst` is included in the sample `autoupg.xml` file. Make sure that the script matches your particular set-up and that it does not remove more than you actually want!

22.5.3 For More Information

- Linuxrc Documentation: <http://en.opensuse.org/SDB:Linuxrc>

Automated Deployment of Preload Images

With KIWI you are able to create operating system images. This chapter describes the process of deploying a system image to an empty client machine. For this, you have to create a preload image which contains a bootable RAW image. This file contains two important parts: a partition table and the actual operating system. This RAW image will be written to the empty hard disk and the operating system extends to the remaining disk space at first boot.

To create such an image, see Section 17.4.2, “Creating an Image” (page 312). When you build the ISO image, you can find the RAW file in the destination folder. There are many ways to dump a raw image onto a disk.

- Plug the disk into a deployment server and just copy the image to the raw device.
- Provide the raw image by means of a HTTP or FTP server and dump it on the disk of the client machine.
- Create a netboot image to get the image and dump it on the disk. This is a good method for mass deployment.
- Boot a rescue disk and do the dump manually from the rescue image.

For a quick start, it is good to use one of the methods described in Section 23.1, “Deploying system manually from rescue image” (page 362).

23.1 Deploying system manually from rescue image

Deploying with generated ISO file from KIWI:

1. Burn the ISO image you get from the KIWI building process, see Section 17.4.2, “Creating an Image” (page 312) on CD/DVD
2. Boot from this medium onto the client machine.
3. Select the hard disk for installation.
4. Restart the client machine and boot from hard disk.

Deploying over rescue system:

1. Boot the client machine with a rescue system. Such systems are available on all SUSE installation CDs or DVDs.
2. Log in as `root`. Do not enter password.
3. Configure your network. If you have DHCP available in your network, this is merely the command `ifup-dhcp eth0`. If you must do this manually, use the command `ip` to configure your network. The output starting DHCP also tells you the IP address of the computer.
4. Listen on an unused port of your network like 1234 and dump the incoming data to disk with the following command:

```
netcat -l -p 1234 > /dev/sda
```
5. On the imaging server, send the raw image to the client machine with the command:

```
netcat <IP of client> 1234 < $HOME/preload_image/<image_name>
```
6. When the image is transferred, remove the rescue system from your CD or DVD drive and shutdown the client machine. On reboot, the boot loader GRUB should be started on the client and the firstboot system will take over.

23.2 Automated Deployment with PXE Boot

When doing multiple installations of an operating system on similar hardware, it is useful to put some effort into preparing a mass deployment of the operating system and to minimize the time needed for the actual deployment. This chapter describes this process. The goal is to simply plug in a computer, connect it to a network, start a network boot, and wait until it powers down.

The following actions have to be performed in order to accomplish this task:

Setup a boot and install server

A dedicated machine is needed, that should be prepared to offer PXE boot as well as an ftp or Web server to provide a preload image. It is a good idea to give the machine enough memory to hold all necessary installation data in memory. For a default installation, you should have at least 4 GByte of memory. All the necessary tasks can be accomplished with SUSE Linux Enterprise Server. For more details, see Section 23.2.1, “Setup a Boot and Install Server” (page 364).

Prepare a preload Image

The actual installation is done by the copying of a raw image of the operating system to the new hard disk. All features and settings must be prepared and tested carefully. To provide such an image, KIWI can be used (available in the SDK of the SUSE Linux Enterprise operating system). More information about image creation with KIWI is available in Chapter 17, *KIWI* (page 307). For more details about the requirements of the preload image, see Section 23.2.2, “Creating a Preload Image” (page 364).

Create an initial system for deployment

This is a task that requires some linux expertise. A description on how this can be achieved by means of an example installation is available at Section 23.2.3, “Creating a Initial System to Deploy a Preload Image” (page 365).

Configure the boot server for automatic deployment

PXE boot must be told to boot the installation system, that in turn will take the preload image from the server and copy it to the hard disk.

23.2.1 Setup a Boot and Install Server

There are four steps to accomplish in order to perform this task after a SUSE Linux Enterprise Server installation:

- 1 Set up the installation source as described in Section 14.2, “Setting Up the Server Holding the Installation Sources” (page 246). Choose an HTTP, or FTP network server.
- 2 Set up a TFTP server to hold a boot image (this image will be created in a later step). This is described in Section 14.3.2, “Setting Up a TFTP Server” (page 258).
- 3 Set up a DHCP server to assign IP addresses to all machines and to reveal the location of the TFTP server to the target system. This is described in Section 14.3.1, “Setting Up a DHCP Server” (page 255).
- 4 Prepare the installation server PXE boot. This is described in further detail in Section 14.3.3, “Using PXE Boot” (page 259).

Note that the actual installation process will greatly benefit if you provide enough memory on this machine to hold the preload image. Also, using gigabit ethernet will speedup the deployment process considerably compared to slower networks.

23.2.2 Creating a Preload Image

The process of creating images with KIWI is described in Section 17.4.2, “Creating an Image” (page 312). However, to create a useful image for mass deployment, several considerations should be taken into account:

- A typical preload image will use the following type:

```
<type primary="true" filesystem="ext3" boot="oemboot/suse-SLES11">vmx</type>
```

- During the setup of a preload image, the image creation process is run multiple times. The repositories needed to build the image should be available on the local computer.
- Depending on the desired usage of the preload, some effort should be invested in configuring firstboot. Find more details about firstboot in Chapter 20, *Deploying*

Customized Preinstallations (page 323). With this method you can also require the user to do initial configurations at the first bootup of the system.

- Many additional features can be configured into the image, like adding update repositories or doing an update on initial bootup. However, it is impossible to describe all possibilities in this document, and (depending on the requirements) the creation of the preload image requires in-depth knowledge of the imaging system KIWI, as well as several other technologies used in SUSE Linux Enterprise Server.

The actual image to be deployed should be available from the ftp or http server that you provided on the installation server.

23.2.3 Creating a Initial System to Deploy a Preload Image

In order to run an automatic deployment, it is necessary to start an initial linux system on the target computer. During a typical installation, the kernel and initial ram file system are read from some boot medium and started by the bios. The needed functionality can be implemented in the ram file system, which together with the kernel will serve as the initial system.

The main features that must be provided by the initial system is the enabling of access to the hard disk and the making available of the network connection. Both of these functions are dependent on the hardware onto which you want to deploy. In theory it is possible to create an initial system from scratch, but to ease this task it is also possible to modify the initial ram file system used by the machine during boot.

The following procedure is just one example of how to create the needed initial ram file system:

- 1 Do a standard installation of SUSE Linux Enterprise Server on the target system.
- 2 Install the package `busybox` on the system.
- 3 Create a new ram file system with the following command:

```
mkinitrd -f busybox -D eth0
```

Note that `eth0` represents the ethernet device to which your network cable is attached. The parameter `-f busybox` adds the multi call binary `busybox`

to the ram file system. After doing this, many standard unix commands are available inside this system.

- 4 Copy the new ram file system and the kernel to your boot server with the command:

```
scp /boot/initrd /boot/vmlinuz pxe.example.com:
```

Replace pxe.example.com with the name of your local boot server or ip address.

- 5 Log into your bootserver as user root, and create a directory where you can modify the ram file system:

```
mkdir ~/bootimage
```

- 6 Change your working directory to this directory with the command `cd ~/bootimage`.

- 7 Unpack the previously copied initial ram file system with the command:

```
zcat ../initrd | cpio -i
```

- 8 Edit the file `run_all.sh`.

- 9 Search for the following line, delete it and the rest of the file:

```
[ "$debug" ] && echo preping 21-nfs.sh
```

- 10 Add the following lines to the end of the files `run_all.sh`:

```
[ "$debug" ] && echo preping 92-install.sh
[ "$debug" ] && echo running 92-install.sh
source boot/92-install.sh
[ "$modules" ] && load_modules
```

- 11 Create a new script `boot/92-install.sh` with the following content:

```
#!/bin/bash
if [ "$(get_param rawimage)" ]; then
    rawimage=$(get_param rawimage)
    if [ "$(get_param rawdevice)" ]; then
        rawdevice=$(get_param rawdevice)
        echo "wget -O ${rawdevice} ${rawimage}"
        wget -O ${rawdevice} ${rawimage}
        sync
        sleep 5
        echo "DONE"
    fi
fi
```

```
fi
# /bin/bash
/bin/poweroff -f
```

12 If you want to have a debug shell before the computer switches off, remove the comment sign before `/bin/bash`.

13 Make this script executable with the command `chmod 755 boot/92-install.sh`.

14 Create a new initial ram file system with the commands:

```
mkdir -p /srv/tftpboot
find . | cpio --quiet -H newc -o | gzip -9 -n > \
/srv/tftpboot/initrd.boot
```

15 Copy the kernel to this directory:

```
cp ../vmlinuz /srv/tftpboot/linux.boot
```

The initial ram file system is now prepared to take two new kernel command line parameters. The parameter `rawimage=<URL>` is used to identify the location of the preload image. Any URL that is understood by `wget` can be used. The parameter `rawdevice=<device>` is used to identify the block device for the hard disk on the target machine.

23.2.4 Boot Server Configuration

The configuration of the boot server is covered in detail in several different chapters as listed in Section 23.2.1, “Setup a Boot and Install Server” (page 364). This section should give a checklist that covers steps that are necessary to configure the system.

- Setup a dhcp server. The subnet where the machines are installed needs the additional lines:

```
filename "pxelinux.0";
next-server 192.168.1.115;
```

In this example, 192.168.1.115 is the ip address of the PXE server `pxe.example.com`.

- Configure a PXE server as described in Section 14.3.3, “Using PXE Boot” (page 259). When editing `/srv/tftpboot/pxelinux.cfg/default`, add the following entries:

```
default bootinstall
label bootinstall
    kernel linux.boot
    append initrd=initrd.boot \
    rawimage=ftp://192.168.1.115/preload/preloadimage.raw rawdevice=/dev/sda
```

- Setup an ftp server and copy your prepared preload image to /srv/ftp/preload/preloadimage.raw.

Test your setup by booting the target system with PXE network boot. This will automatically copy the prepared preload image to hard disk and switch off the machine when ready.



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 noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

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

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

1. APPLICABILITY AND DEFINITIONS

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

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

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

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

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

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

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

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

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

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

2. VERBATIM COPYING

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

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

3. COPYING IN QUANTITY

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

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

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

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

4. MODIFICATIONS

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

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION

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

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

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

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

ADDENDUM: How to use this License for your documents

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

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

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

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

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

