SuSE

SuSE Linux

Enterprise Server 8

for IBM S/390 and IBM zSeries

Installation

Contents

Introduction

About This Document

This book describes the procedure for installing SuSE Linux Enterprise Server on IBM S/390 and zSeries systems. You will be provided with all information needed to prepare an installation on the OS/390 and z/OS side as well as the installation procedure of SuSE Linux Enterprise Server itself.

Wherever possible, links are provided to more specific and most up-to-date documentation both on the net and in your installed Linux system. Detailed information about Linux device drivers is excellently covered by IBM (online) documentation.

Structure

Basically, this installation guide is divided in two parts.

Basic Preparations   This part provides information about the requirements both on the hardware and the software side that need to be met to install SuSE Linux Enterprise Server successfully on your machine. Learn out which basic preparations need to be performed on the OS/390 and z/OS side.

Installing SuSE Linux Enterprise Server on IBM S/390 and zSeries   This part covers the complete installation procedure including the preparation of the installation system, network setup, and a complete in-depth description of the YaST2 install procedure. A summary of all supported network connection types and some short examples are also provided.

Additionally, the appendix contains five chapters about automated configuration using the parmfile, support of fibre channel attached SCSI devices, handling YaST2 in text mode, LVM (Logical Volume Manager), and an overview of Linux file systems.

Required Background

Several assumptions have been made concerning your background knowledge when writing this document.

- You are familiar with S/390 and zSeries terminology.

- You have a good knowledge of the S/390 and zSeries devices attached to your system, especially its network environment.

- You have a basic understanding of handling a Linux or Unix system.

Typographical Conventions

The following typographic conventions are used in this book:

Example	Meaning
YaST	programs
/etc/passwd	files or directories
⟨*placeholder*⟩	replace the character string `placeholder` (including the angle brackets) with the actual value
PATH	an environment variable called PATH
192.168.1.2	the value of a variable
ls	commands
user	users
earth:~ # **ls**	enter the command `ls` in the shell of the user `root` in his home directory on the host "Earth"

Table 1: *continued overleaf...*

```
newbie@earth:~ >
    ls                      enter the command ls in the shell of the user
                            newbie in his home directory on the host
                            "Earth"
C:\> fdisk                  at the DOS prompt, enter the command fdisk
(Alt)                       press this key; keys to press sequentially are
                            separated by spaces
(Ctrl) + (Alt) + (Del)      keys to press simultaneously are connected with
                            a '+'
"Permission denied"         system messages
'System update'             menu item or button text 'System update'
```

Acknowledgments

The history of Linux is a success story about countless developers all around the world contributing to what originally started as a "one-man show" by Linus Torvalds. Thanks to all of them for their tremendous efforts.

Especially we would like to thank all the people involved in the Linux for zSeries and S/390 project at IBM and SuSE. Thanks to:

- the developers at SuSE and IBM
- the test team at SuSE
- all beta-testers and proofreaders at IBM

Thanks for making SuSE Linux Enterprise Server for S/390 and zSeries possible.

Nuremberg, 26th January 2004

      Your SuSE team

Part I

Basic Preparations

System Requirements

This chapter gives basic information about the system requirements. At first, a list of hardware running SuSE Linux Enterprise Server for S/390 and zSeries is provided. The level of the MicroCode (MCL) used in your S/390 or zSeries system is very important for the installation. Additional software to install and use for installation is mentioned at the end of this chapter.

Hardware

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

S/390

- IBM S/390 Multiprise 3000 Enterprise Server (MP3000) 7060
- IBM Parallel Enterprise Server G5 (G5) 9672
- IBM Parallel Enterprise Server G6 (G6) 9672

zSeries

- IBM eServer zSeries 800 (z800) 2066
- IBM eServer zSeries 900 (z900) 2064
- IBM eServer zSeries 990 (z990) 2084

Memory Requirements

Different installation methods have different memory requirements during installation. After installation is completed, the system adminstrator may reduce memory to the desired size.

SuSE recommends using:

128+ MB	for installation using ssh (text mode) via NFS or SMB accessible installation sources
256+ MB	for installation using X11 (graphical mode) from NFS or SMB installation sources
512 MB	for installation using VNC (graphical mode remotely displayed in Java-enabled web browser) via FTP accessible installation sources. The increased memory requirement during installation results from using temporary disk space in memory (RAM disk) during the package transfers.

DASD Requirements

The following DASD sizes are required for the installation of SuSE Linux Enterprise Server:

- Default installation 31-bit: 1.3 GB

- Default installation 64-bit: 1.6 GB

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-2

- OSA Express Ethernet (including Fast and Gigabit Ethernet)

- Token Ring

- Escon

- CTC (or virtual CTC)

- IUCV

- Hipersockets/Guest LAN

- OSA Express High-Speed Token Ring

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 ramdisk). If this option is not available and you cannot use z/VM for installing the system, you need to IPL from a tape with the tapeipl kernel, the parmfile, and the initrd. Thus, you need access to a tape unit (3480, 3490, or 3590, for example).

MicroCode Level and APARs/Fixes

This release of SuSE Linux Enterprise Server is based upon the June 2003 codestream from IBM Developerworks (http://www10.software.ibm. com/developerworks/opensource/linux390/june2003_technical. shtml). The restrictions and requirements listed at the above website do also apply for this release of SuSE Linux Enterprise Server, except where stated differently in this manual.

For installation under VM, you need at least z/VM 4.2 with service level 0103. To use Hipersockets under VM on 2064/z900, you need MicroCode EC E26949 level 013 or higher.

For the installation of SuSE Linux Enterprise Server 8 on IBM S/390 or zSeries, the following MicroCode Levels and z/VM APARs are required:

OSA-Express QDIO

zSeries 990 MSL

zSeries 900 GA3
  Driver 3G, OSA microcode level 3.0A
  MCLs: J11204.007 and J11204.008 (available May 03, 2002)

zSeries 900 GA2
  Driver 3C, OSA microcode level: 2.26
  MCLs: J10630.013 and J10630.014 (available May 20, 2002)

zSeries 800 GA1
  Driver 3G, OSA microcode level 3.0A
  MCLs: J11204.007 and J11204.008 (available May 03, 2002)

S/390 Parallel Enterprise Servers G5 and G6
  Driver 26, OSA microcode level: 4.25
  MCLs: F99904.032 and F99904.033 (available May 16, 2002)

VM/ESA and z/VM

- **z/VM 4.4**
  All necessary fixes and enhancements included.

- **z/VM 4.3**
  All necessary fixes and enhancements included.

- **z/VM 4.2**
  APAR: VM63181, PTF: UM30610
  APAR: VM62938, PTF: UM30225
  APAR: VM63034, PTF: UM30290

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.

Software

To install SuSE Linux Enterprise Server via non-Linux based NFS or FTP, you might experience problems with NFS/FTP server software. Especially Windows standard FTP server can cause errors, so we generally recommend installing via SMB on these machines.

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

- **Telnet or ssh with terminal emulation (xterm compatible)**
  ssh and telnet are standard Unix tools that should be present on any Unix or Linux system.

  For Windows, there is a telnet and ssh client called Putty. It is free to use and is included on CD 1 in the directory `/dosutils/putty`. More information about Putty can be obtained at:
  http://www.chiark.greenend.org.uk/~sgtatham/putty.html.
  The use of Putty is recommended because the standard telnet implementation for Windows is not fully compliant to Unix standards and not sufficient for Linux applications.

- **VNC client**
  For Linux, a VNC client called vncviewer is included in SuSE Linux as part of the `vnc` package.

  For Windows, a VNC client is included in the present SuSE Linux Enterprise Server. Find it in `/dosutils/vnc-3.3.3r9_x86_win32`.

`tgz` of CD 1. After copying this file to your hard disk, unpack it using untgz, which can be found in `/dosutils/untgz` on CD 1. The VNC client vncviewer.exe will then reside under `vnc_x86_win32\vncviewer`.

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 MI/X (MicroImages X Server) can be obtained at:
  `http://www.microimages.com/mix`

Tip

Additional Information

Consult the README located in the highest directory of CD 1 of your SuSE Linux Enterprise Server prior to installing SuSE Linux Enterprise Server on S/390 or zSeries. This file completes the documentation presented in this book.

Tip

Types of Installation

This chapter gives an overview of the different types of installation possible with SuSE Linux Enterprise Server S/390 and zSeries. Basically, these three types are given:

- Native
  Installation of SuSE Linux Enterprise Server as a native operating system on your S/390 or zSeries

- LPAR
  Installation of SuSE Linux Enterprise Server using a logical partition (LPAR)

- VM (z/VM)
  Installation of SuSE Linux Enterprise Server as an operating system within z/VM

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

Native

Installing SuSE Linux Enterprise Server for S/390 and zSeries in native mode means that SuSE Linux Enterprise Server will be the only operating system running on your hardware. SuSE Linux Enterprise Server uses all physical memory and all processors of your zSeries. You cannot run other operating systems simultaneously on your system.

LPAR

If you install SuSE Linux Enterprise Server S/390 and zSeries into a separate logical partition (LPAR), you allow SuSE Linux Enterprise Server to use a special part of physical memory in your system. You can also decide how many processors are used by SuSE Linux Enterprise Server. In this mode, you are able to run different operating systems simultaneously on your S/390 or zSeries system.

z/VM

Running SuSE Linux Enterprise Server for S/390 and zSeries 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.

IPL Options

This chapter 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 options Tape, VM reader, and Load From CD-ROM or Server are discussed. Installing the software packages, which is done via the network, does not require the IPL medium.

Tape

IPLing from 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>`.

VM Reader

To IPL from 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 comfortable 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.

Load From CD-ROM or Server

For IPLing into an LPAR, there is the option to directly load from the HMC CD-ROM device. It also allows you to IPL via FTP from a server holding the SuSE Linux Enterprise Server zSeries files. Find this button near the 'LOAD' button (if you have upgraded to the latest MCL, see Chapter *System Requirements* on page 7).

This function can be performed from the HMC. This is done by accessing the 'Defined CPCs' icon in the 'Groups Work Area' on the HMC desktop. From there, access the 'Defined CPCs Work Area'. Select a system image and initiate 'CPC Recovery'. Use 'Single Object Operations' to remotely access the Service Element.

Access the Service Element's 'Groups Work Area'. Select the LPAR to perform the load from CD-ROM or server and initiate 'CPC recovery' for that LPAR. Initiate the 'Load from CD ROM or Server' from the action menu. Provide information as requested on subsequent panels.

The IOCDS

This chapter provives some necessary information about the IOCDS and how to customize some settings for sharing network cards or DASD among several LPARs. In the IOCDS, the chpid and types of the devices connected to the S/390 or zSeries are defined. The resources can be shared or dedicated between LPARs.

┌─ **Caution** ───

Sharing Devices (DASD)
Do not share DASD writable between LPARs as 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 S/390 or zSeries.

─────────────────────────────────────── **Caution** ─┘

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

```
CHPID PATH=FD,TYPE=DSD,SHARED
CNTLUNIT CUNUMBR=FD00,PATH=FD,UNITADD=((00,256)),UNIT=3990-2
IODEVICE ADDRESS=(FD03,1),CUNUMBR=FD00,UNIT=3390,PART=LPAR1
```

Output 1: Dedicating DASD to One LPAR

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

Several Linux systems can use the same network device if you share it between 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. These modes can be defined 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 S/390 or zSeries Linux environment.

```
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
```

Output 2: Sharing OSA Express Card between LPARs (non-qdio)

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

Output 3: Sharing OSA Express Card between LPARs (qdio)

Making the Installation Data Available

This chapter provide detailed informations about making the SuSE Linux Enterprise Server S/390 and zSeries 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 S/390 and zSeries system.

> **Tip**
>
> To IPL from CD-ROM, it is reasonable to make a copy of the CD or copy its contents onto the FTP, NFS, or SMB server first.
>
> **Tip**

Creating an IPLable Tape

If you cannot IPL from CD-ROM or from VM reader you'll need to create a tape from which to IPL the SuSE Linux Enterprise Server installation image. There are many ways to achieve this. One is to copy the files:

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

as binary files from CD 1 into the z/VM (for example by FTP from a Linux workstation).

Name them

```
SLES8 IMAGE
SLES8 PARM
SLES8 INITRD
```

and write them onto a tape with the REXX script seen in Output 4.

Note ───

Transferring Binaries using FTP

Do not upload the files as `fixed 80`. Store them as `fixed 1024`. Use
the ftp command `loc site f 1024`.

── **Note** ─┘

```
'REWIND 181'
'FILEDEF IN1 DISK' SLES8 IMAGE A
'FILEDEF IN2 DISK' SLES8 PARM A
'FILEDEF IN3 DISK' SLES8 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
```

Output 4: REXX Script to Create an IPLable Tape

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

Using a Linux Workstation or SuSE Linux CD

If you have a Linux workstation running in your computer environment, use
the workstation to provide the installation data to the S/390 and zSeries in-
stallation process by NFS or FTP.

Via NFS

Use NFS (Network File System) for making the installation media available.

> **Note**
>
> **Exporting Mounted Devices via NFS**
>
> Exporting the file system root (/) does not imply the export of mounted devices, such as CD-ROM. Explicitly name the mount point in /etc/exports:
>
> /media/cdrom *(ro)
>
> After changing this file, restart the NFS server. On SuSE Linux Enterprise Server, restart the NFS server simply by typing rcnfsserver restart.
>
> **Note**

Via FTP

Setting up an FTP server on a Linux system involves the installation of the server software itself (e. g., wuftpd or proftpd) as well as other possible configuration tasks. Using YaST2, 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 FTP login with a valid user name and password. You might want to create a user account for this task only. The FTP daemon itself 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

```
earth:/root #   rcinetd restart
```

SuSE Linux on CD

CD 1 of the SuSE Linux Enterprise Server S/390 and zSeries contains a bootable Linux image for Intel-based workstations. Boot from this CD, 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.

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. Details on this can be found 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.

Via SMB

To make the installation media available via SMB, just insert the SuSE Linux Enterprise Server CD 1 into the CD drive of the Windows workstation. Then create a new share using the CD-ROM drive's letter and make it available for everyone in the network.

> **Tip**
>
> **SMB Installation Path for YaST2**
> Consult the README from the highest directory of CD 1 for detailed information about the syntax of the installation path for YaST2.
>
> **Tip**

Via NFS

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

Via FTP

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

The FTP server that is part of 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 Humingbird Exceed or WAR-FTPD, have been reported as functional.

Using VNC for Installation

VNC "Virtual Network Computing" is a client server solution to control a remote X server via a small and simple to use client. This client is available for many operating systems, including various Microsoft Windows versions, Apple's MacOS, and Linux.

Uses for VNC

SuSE Linux Enterprise Server uses VNC to enable the use of the graphical user interface of the installation software, YaST2. It is needed for systems that cannot provide a graphical console or if the administrator cannot access the graphical console (e.g., a rack mounted system without a monitor).

The VNC client vncviewer is used to enable the display and control of YaST2 during the installation process. Before booting for the installation, prepare a remote computer, which is connected via network to the system on which to install.

Preparing a SuSE Linux Client

SuSE Linux Enterprise Server can also be used as a control client for the installation. The VNC client vncviewer is part of the package vnc. Use the YaST2 'Install/Remove Software' module to install the package.

Figure 6.1: YaST2 on a SuSE Linux Desktop

Preparing a Microsoft Windows Client

The VNC client for all Microsoft Windows 32-bit versions can be found in the
dosutils/vnc directory on CD 1. Copy the file vnc-3.3.3r9_x86_win32.tgz
to a directory on your hard disk and unpack the file with the untgz utility,
which can be found in the dosutils/untgz directory on CD 1.

```
C:\ tmp > untgz vnc-3.3.3r9_x86_win32.tgz
```

The VNC client vncviewer.exe can be found in the
vnc_x86_win32\vncviewer directory.

Alternatively, install SuSE Linux Enterprise Server using the VNC Java client.

Part II

Installing on IBM
S/390 and zSeries

Installation Types on S/390 and zSeries

As described in Section *Types of Installation* on page 13, there are three differ-ent installation modes for Linux on S/390 and zSeries:

- ESA Native Installation

- LPAR Installation

- z/VM Installation

This chapter provides information about which steps must be performed to install SuSE Linux Enterprise Server Enterprise Server for each of the installa-tion modes and where to find the appropriate information. After the prepara-tions mentioned in the previous chapters have been accomplished, follow the installation overview of the desired installation mode below to install SuSE Linux Enterprise Server on your system.

Note

As ESA/Native and LPAR installation closely resemble each other, some sections describe the procedure for both of them. Installation with YaST2 is summarized in one chapter. Minor differences between the installation types are appropriately highlighted.

Note

Installation Overview of ESA Native Installation

1. Prepare the devices needed for installation. See Chapter *Preparing the IPL of the SuSE Linux Installation System* Section *ESA Native Installation and LPAR Installation* on page 33.

2. IPL the installation system. See Chapter *IPLing the SuSE Linux Installation System* Section *ESA Native Installation* on page 39.

3. Configure the network. See Chapter *Network Configuration* Section *ESA Native Installation and LPAR Installation* on page 45.

4. Connect to the SuSE Linux Enterprise Server installation system. See Chapter *Connecting to the SuSE Linux Installation System*.

5. Start installation using YaST2. See Chapter *Installation* on page 61.

6. IPL the installed system. See Chapter *IPLing the Installed System* Section *ESA Native Installation and LPAR Installation* on page 73.

Installation Overwiew of LPAR Installation

1. Prepare the devices needed for installation. See Chapter *Preparing the IPL of the SuSE Linux Installation System* Section *ESA Native Installation and LPAR Installation* on page 33.

2. IPL the installation system. See Chapter *IPLing the SuSE Linux Installation System* Section *LPAR Installation* on page 40.

3. Configure the network. See Chapter *Network Configuration* Section *ESA Native Installation and LPAR Installation* on page 45.

4. Connect to the SuSE Linux Enterprise Server installation system. See Chapter *Connecting to the SuSE Linux Installation System*.

5. Start installation using YaST2. See Chapter *Installation* on page 61.

6. IPL the installed system. See Chapter *IPLing the Installed System* Section *ESA Native Installation and LPAR Installation* on page 73.

Installation Overwiew of z/VM Installation

1. Prepare the devices needed for installation. See Chapter *Preparing the IPL of the SuSE Linux Installation System* Section *z/VM Installation* on page 34.

2. IPL the installation system. See Chapter *IPLing the SuSE Linux Installation System* Section *z/VM Installation* on page 41.

3. Configure the network. See Chapter *Network Configuration* Section *z/VM Installation* on page 51.

4. Connect to the SuSE Linux Enterprise Server installation system. See Chapter *Connecting to the SuSE Linux Installation System*.

5. Start installation using YaST2. See Chapter *Installation* on page 61.

6. IPL the installed system. See Chapter *IPLing the Installed System* Section *z/VM Installation* on page 73.

Preparing the IPL of the SuSE Linux Installation System

ESA Native Installation and LPAR Installation

Note

Configuration for ESA Native Installation
You need to configure your zSeries system to start in ESA mode with an appropriate activation profile and IOCDS. Consult IBM documentation about how to achieve this.

Note

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). Output 5 shows an example that 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 IBM documentation (http://www.s390.ibm.com/linux/press.html).

```
CHPID PATH=FD,TYPE=DSD,PART=(Z1)
CHPID PATH=FE,TYPE=OSA,PART=(Z1)
CHPID PATH=10,TYPE=CNC,PART=(Z1)

CNTLUNIT CUNUMBR=FD00,PATH=FD,UNITADD=((00,1)),UNIT=3990-2
IODEVICE ADDRESS=(FD00,1),CUNUMBR=FD00,UNIT=3390

CNTLUNIT CUNUMBR=FE20,PATH=FE,UNIT=OSA,UNITADD=((20,1))
```

```
IODEVICE ADDRESS=(FE20,1),CUNUMBR=FE20,UNIT=OSA,PART=Z1

CNTLUNIT CUNUMBR=100A,PATH=10,UNIT=3480,UNITADD=((0A,1))
IODEVICE ADDRESS=100A,CUNUMBR=100A,UNIT=3480,PART=Z1
```

Output 5: An Example IOCDS

Where Next

For ESA Native Installation proceed with Section *ESA Native Installation* on page 39 in Chapter *IPLing the SuSE Linux Installation System*. For LPAR Installation, proceed with Section *LPAR Installation* on page 40 in Chapter *IPLing the SuSE Linux Installation System*.

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. Output 6 shows the directory for a user "LINUX1" with the password "LIN-PWD", 128 MB of memory (extendable up to 256 MB), 32 MB of expanded RAM (XSTORE) and some minidisks (MDISK), two CPUs, an OSA QDIO device, and virtual CTC definitions.

Tip

Assigning Memory to z/VM guests

When assigning memory to a z/VM guest, make sure the memory size suits the needs of your preferred installation type. For more information, refer to Section *Memory Requirements* on page 7.

Tip

```
USER LINUX1 LINPWD 128M 256M G
*-------------------------------------------
* LINUX1
*-------------------------------------------
* This VM Linux guest has two CPs defined.

    CPU 01 CPUID 111111
    CPU 02 CPUID 111222
```

```
IPL CMS PARM AUTOCR
IUCV ANY
IUCV ALLOW
MACH ESA 10
OPTION MAINTCCW RMCHINFO
SHARE RELATIVE 2000
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 0254 2800 DASD40 MR ONE4ME TWO4ME THR4ME
```

Output 6: Configuration of a z/VM Directory

We will use minidisk 201 as the guest's home disk. Minidisk 150 with 200 cylinders will be the Linux swap device. Disk 151 with 3000 cylinders will hold the Linux installation.

Now add (as the user "MAINT") the guest to the user directory:

```
DIRM FOR LINUX1 ADD
```

Enter the name of the guest ("LINUX1") and press F5. Set up the environment of the user now 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)
```

Now assign the directories to the guest:

```
DIRMAP USER DIRECT A
```

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

Setting up CTC

z/VM guests often communicate through CTC connections with each other. This requires some setup on the z/VM host and the z/VM guest.

The virtual CTC devices were defined in Output 6 on page 34. Now, couple the virtual CTC devices in LINUX with the virtual CTC devices in TCPIP.

Add the following lines to the user's (i.e., "LINUX1") PROFILE EXEC and adjust them to your needs:

```
DEF CTC 120
DEF CTC 121
COUPLE 120 TCPIP 01F1
COUPLE 121 TCPIP 01F0
```

In this example, the CTC port pair will have the numbers 120 and 121 and access port 1F1 and 1F0 on the z/VM host.

> **Note**
>
> **Setting up CTC**
> Make sure to couple an even (read) address with an odd (write) address.
>
> **Note**

On the z/VM host, first make sure to have CTC devices defined for user TCPIP in USERDIRECT:

```
SPECIAL 01F0 CTCA
SPECIAL 01F1 CTCA
```

Now configure these ports as user TCPMAINT by changing the file PROFILE TCPIP. First add the appropriate DEVICE and LINK statements, in our example:

```
DEVICE LIN1CTC CTC 1F0
LINK LIN1C CTC 0 LIN1CTC
```

Assign the IP address for the z/VM host to this interface in the HOME section now:

```
HOME
   192.168.0.17 LIN1C
```

Then add a static route to the z/VM guest in the GATEWAY section

```
GATEWAY
  192.168.0.18 = LIN1C 1500 HOST
```

Finally, start the interface with the following line at the end of PROFILETCPIP:

```
START LIN1CTC
```

Save your changes, restart the TCP/IP subsystem with

```
OBEYFILE PROFILE TCPIP B
```

In case this fails although you have made sure that PROFILETCPIP is correct, try to restart the TCP/IP subsystem as user "MAINT":

```
FORCE TCPIP
XAUTOLOG TCPIP
```

Caution

Restarting the TCP/IP Subsystem

FORCE TCPIP disconnects all TCP/IP clients from the service, making them unreachable through TCP/IP. This includes running VM guests. Keep VM usable by running the above commands from a terminal that is not connected over TCP/IP (e.g., a directly-connected 3270 terminal or the HMC) and shutting down all VM guests and services using TCP/IP prior to the execution of FORCE TCPIP. This applies to Linux guests using CTC or IUCV through the TCPIP user.

Caution

Where Next

Proceed with Section *z/VM Installation* on page 41 in Chapter *IPLing the SuSE Linux Installation System*.

IPLing the SuSE Linux Installation System

ESA Native Installation

There are different ways to IPL SuSE Linux natively on the S/390 and zSeries system. The preferred way is the "Load from CD-ROM or server" feature of the SE or HMC.

IPL from CD-ROM

Mark the 'Images' icon 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 CD-ROM then 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 Tape

Prepare a tape with the files `tapeipl.ikr`, `parmfile`, and `initrd`. The files must be on the tape in this order. Chapter *Making the Installation Data Available* on page 19 describes an example of how to create the tape 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.

Emulated I/O

This option is only available on the Multiprise architecture. Insert the SuSE Linux Enterprise Server CD into the CD-ROM and select 'LOAD' with the device address for Emulated I/O as set up in the IOCDS. If the LOAD fails with a busy error, simply try again a few times. Do a "rewind" of the CD-ROM prior to each LOAD (remember that it is emulated as a tape) by entering

```
F:\> awsmount 080 /REW /D /R
```

on the OS/2 console. `080` is the device number of the emulated CD-ROM. There will be a fix for that in the near future provided by IBM. The "Operating system messages" should now show the kernel boot messages.

Where Next

Proceed with Section *ESA Native Installation and LPAR Installation* on page 45 in Chapter *Network Configuration*.

LPAR Installation

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

IPL from CD-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 CD-ROM then 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 Tape

Follow the procedure described for ESA Native Installation (see Section *ESA Native Installation* on the page before).

Emulated I/O

Follow the procedure described for ESA Native Installation (see Section *ESA Native Installation* on page 39).

Where Next

Proceed with Section *ESA Native Installation and LPAR Installation* on page 45 in Chapter *Network Configuration*.

z/VM Installation

This chapter is about IPLing the installation system to install SuSE Linux Enterprise Server S/390 and zSeries on a z/VM system.

Transfer of the Installation System via FTP

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

Login as the z/VM Linux guest to IPL. Make the content of the directory /boot/ on CD 1 of the SuSE Linux Enterprise Server Enterprise Server for S/390 and zSeries available by FTP within your network. From this directory, get the files vmrdr.ikr, initrd, and parmfile. Output 7 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 will be different for your network.

```
FTP 192.168.0.3
VM TCP/IP FTP Level 320
Connecting to 192.168.0.3, port 21
220 ftpserver FTP server (Version wu-2.4.2-academ[BETA-18](1)
Wed Mar 28 19:16:59 GMT 2001) 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/cdrom/boot/vmrdr.ikr sles8.image
200 PORT Command successful
150 Opening BINARY mode data connection for vmrdr.ikr
(1581060 bytes)
226 Transfer complete.
1581060 bytes transferred in 2.065 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
get /media/cdrom/boot/initrd sles8.initrd
200 PORT Command successful
150 Opening BINARY mode data connection for initrd
(12194534 bytes)
226 Transfer complete.
12194534 bytes transferred in 15.919 seconds.
Transfer rate 766.70 Kbytes/sec.
Command:
asc
200 Type set to A
Command:
get /media/cdrom/boot/parmfile sles8.parm
150 Opening BINARY mode data connection for parmfile
(38 bytes)
226 Transfer complete.
38 bytes transferred in 0.092 seconds.
Transfer rate 0.38 Kbytes/sec.
Command:
quit
```

Output 7: *Transferring the Binaries via FTP*

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 that must not get
converted, so use the "binary" transfer mode.

Creating a Start Script

It is convenient to write a small REXX script to IPL the Linux installation sys-
tem. This script loads the kernel, parmfile, and the initial ramdisk into the

reader for IPL (Output 8).

```
        XEDIT SLES8 EXEC

/**/
'close rdr'
'purge rdr all'
'spool punch * rdr'
'PUNCH SLES8 IMAGE A (NOH'
'PUNCH SLES8 PARM A (NOH'
'PUNCH SLES8 INITRD A (NOH'
'change rdr all keep nohold'
'ipl 00c clear'
```

Output 8: SLES8 EXEC

Save the script with the command FILE.

Now you can IPL the SuSE Linux Enterprise Server Enterprise Server Installation System with the command:

```
        SLES8
```

The Linux kernel then starts and prints its boot messages.

Where Next

Proceed with Section *z/VM Installation* on page 51 in Chapter *Network Configuration*.

Network Configuration

ESA Native Installation and LPAR Installation

After the kernel has completed its start-up routines (Output 9), open the "Operating System Messages" on the HMC or SE. There, answer a few questions regarding the network setup.

```
Linux version 2.4.19-2suse-SMP (root@hostname) (gcc ver-
sion 3.2) #1 SMP Wed Sep 18 13:08:12 UTC 2002
We are running native
This machine has an IEEE fpu
On node 0 totalpages: 131072
zone(0): 131072 pages.
zone(1): 0 pages.
zone(2): 0 pages.
Building zonelist for node : 0
Kernel command line: ramdisk_size=65536 root=/dev/ram1 ro cio_msg=no
Highest subchannel number detected (hex) : 0012
Calibrating delay loop... 788.07 BogoMIPS
Memory: 507608k/524288k available (1552k kernel code, 0k re-
served, 389k data, 56k init)
Dentry cache hash table entries: 65536 (order: 7, 524288 bytes)
Inode cache hash table entries: 32768 (order: 6, 262144 bytes)
Mount-cache hash table entries: 512 (order: 0, 4096 bytes)
Buffer-cache hash table entries: 32768 (order: 5, 131072 bytes)
Page-cache hash table entries: 131072 (order: 7, 524288 bytes)
debug: Initialization complete
POSIX conformance testing by UNIFIX
Detected 2 CPU's
Boot cpu address  1
cpu 0 phys_idx=1 vers=FF ident=111111 machine=2064 unused=0000
cpu 1 phys_idx=2 vers=FF ident=111222 machine=2064 unused=0000
migration_task 0 on cpu=0
migration_task 1 on cpu=1
```

```
init_mach : starting machine check handler
Linux NET4.0 for Linux 2.4
Based upon Swansea University Computer Society NET3.039
Initializing RT netlink socket
mach_handler : ready
mach_handler : waiting for wakeup
Starting kswapd
kinoded started
aio_setup: num_physpages = 32768
aio_setup: sizeof(struct page) = 44
pty: 256 Unix98 ptys configured
RAMDISK driver initialized: 16 RAM disks of 65536K size 1024 blocksize
loop: loaded (max 16 devices)
md: md driver 0.90.0 MAX_MD_DEVS=256, MD_SB_DISKS=27
md: Autodetecting RAID arrays.
md: autorun ...
md: ... autorun DONE.
debug: cio_msg: new level 6
debug: cio_trace: new level 6
debug: cio_crw: new level 6
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
IP: routing cache hash table of 2048 buckets, 32Kbytes
TCP: Hash tables configured (established 65536 bind 65536)
Linux IP multicast router 0.06 plus PIM-SM
NET4: Unix domain sockets 1.0/SMP for Linux NET4.0.
RAMDISK: Compressed image found at block 0
Freeing initrd memory: 7181k freed
VFS: Mounted root (ext2 filesystem).
done doing movetotmpfs
IPv6 v0.8 for NET4.0
IPv6 over IPv4 tunneling driver
Mounted /proc
Creating /var/log/boot.msg
Enabling system logging...
Sep 19 08:00:05 suse syslogd 1.4.1: restart.
```

Output 9: Kernel Messages during the Boot Sequence

First, select the type of network connection to use: Token Ring, Ethernet, Express Gigabit Ethernet, CTC, ESCON, IUCV, Hipersockets, or Express High-speed Token Ring (Output 10). In this example installation, ethernet is used, so 2 is selected. Enter the device address of the ethernet network device next (FC20 in our example). Then enter the port to use. SuSE Linux Enterprise Server now tries to load the network module by building a parameter line with the information provided then displays all loaded modules. Loading was successful a module named lcs is listed.

```
=                                                         =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                         =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA-Express High-Speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 10: Supported Network Connection Types

When prompted, enter the network device number of your network card
(e.g., FC20) and the relative port on this device (Output 11).

```
Enter the read channel device number, e.g. 'FC20' (0c30):
fc20
Please enter the relative port number on device address fc20
Relative port, e.g. '0' (0):
2
Writing 'noauto;lcs0,0xfc20,0xfc21,0,1' to /proc/chandev
Using /lib/modules/2.4.version/net/lcs.o
Starting lcs module
with chandev support,with multicast support, with ethernet support,
with token ring support.
debug: lcs: new level 0
eth0: ip v6 supported yes enabled yes
eth0: multicast supported yes enabled yes
lcs: eth0 configured as follows read subchannel=162 write
subchannel=163
read_devno=fc20 write_devno=fc21
hw_address=00:04:AC:E3:53:88 rel_adapter_no=1
lcs                    27584   0 (unused)
eth0      Link encap:Ethernet   HWaddr 00:04:AC:E3:53:88
BROADCAST MULTICAST  MTU:1500  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:100
RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
eth0 is available, continuing with network setup.
```

Output 11: *Network Device Driver Parameters*

For information about the driver parameters, read the manual "LINUX
for zSeries Device Drivers and Installation Commands for Kernel 2.4" at
http://www10.software.ibm.com/developerworks/opensource/
linux390/june2003_documentation.shtml.

Next, enter the full host name, the IP address, and the netmask (Output 12).
If you install via Token Ring, Ethernet, Gigabit Ethernet, or High-Speed To-
ken Ring, also enter the broadcast address and the gateway address. For
CTC, ESCON, and IUCV, the gateway is the IP address of your peer. Finally,
the IP address of the DNS server, the DNS search domain, and the MTU are
requested. The MTU size should always match the one used by the network
to which to connect.

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.20
Please enter the net mask:
 255.255.255.0
Please enter the broadcast address:
 192.168.0.255
Please enter the gateway's IP address:
 192.168.0.1
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain:
 cosmos.com
Please enter the MTU (Maximum Transfer Unit),
leave blank for default: (1500):
```

Output 12: *Network Device Name*

A summary is then displayed. Confirm whether your input is correct (Out-
put 13). 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. See Output 14 on the facing page).

```
Configuration for eth0 will be:
Full host name   : earth.cosmos.com
IP address       : 192.168.0.20
Net mask         : 255.255.255.0
Broadcast address: 192.168.0.255
Gateway address  : 192.168.0.1
DNS IP address   : 192.168.0.2
DNS search domain: cosmos.com
MTU size         : 1500
Is this correct (Yes/No) ? y
```

Output 13: *Network Settings Summary*

```
For security reasons you have to set an temporary installation
system password for the user "root".
You'll be asked for it only when you telnet into this installation
system to limit the access to it and it will be cleared as soon
as you shut down or reset the installation system

Please enter the temporary installation password: suse
```

Output 14: *Entering the Password*

With all basic parameters, set up the network to start. Take a look at the output of ifconfig, which should contain two entries: a loopback (lo) connection and one connection out of tr0, eth0, ctc0, escon0, iucv0 or hsi0 with correct settings (Output 15). To ensure that all IP addresses are reachable, a ping is sent to the local IP, the gateway, and the name server.

```
Temporary installation password set.
ifconfig eth0 192.168.0.20 netmask 255.255.255.0
broadcast 192.168.0.255 mtu 1500
debug: lcs: new level 0
/sbin/ifconfig eth0 :
eth0      Link encap:Ethernet  HWaddr 00:02:55:09:D2:5D
          inet addr:192.168.0.20  Mask:255.255.255.0
          UP RUNNING NOARP MULTICAST  MTU:1500  Metric:1
          RX packets:2 errors:0 dropped:0 overruns:0 frame:0
          TX packets:3 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:56 (56.0 b)  TX bytes:348 (348.0 b)
          Interrupt:11
Trying to ping my IP address:
PING 192.168.0.20 (192.168.0.20): 56 data bytes
```

```
64 bytes from 192.168.0.20: icmp_seq=0 ttl=255 time=1.625 ms
64 bytes from 192.168.0.20: icmp_seq=1 ttl=255 time=0.109 ms
64 bytes from 192.168.0.20: icmp_seq=2 ttl=255 time=0.062 ms
 --- 192.168.0.20 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.062/0.598/1.625 ms
Waiting 12s for network device eth0.............done.
Trying to ping the IP address of the Gateway:
PING 192.168.0.1 (192.168.0.1): 56 data bytes
64 bytes from 192.168.0.1: icmp_seq=0 ttl=255 time=0.997 ms
64 bytes from 192.168.0.1: icmp_seq=1 ttl=255 time=0.566 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=255 time=0.621 ms
 --- 192.168.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.566/0.728/0.997 ms
Gateway seems to respond to our pings, continuing.
Trying to ping the IP address of the DNS Server:
PING 192.168.0.2 (192.168.0.2): 56 data bytes
64 bytes from 192.168.0.2: icmp_seq=0 ttl=255 time=1.215 ms
64 bytes from 192.168.0.2: icmp_seq=1 ttl=255 time=0.675 ms
64 bytes from 192.168.0.2: icmp_seq=2 ttl=255 time=0.707 ms
 --- 192.168.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.675/0.865/1.215 ms
Network Setup finished, running inetd...

You should be able to login via telnet now, for ssh wait a
few seconds, temporary host keys (only for installation) are
being generated now:
...
Starting SSH daemon..done
...
Generation of temporary installation host keys finished.
After installation, new, different SSH keys will be generated.
```

Output 15: Example ifconfig

Where Next

Proceed with Chapter *Connecting to the SuSE Linux Installation System.*

z/VM Installation

After the kernel has completed its start-up routines (Output 16), answer a few questions regarding to the network setup.

```
Linux version 2.4.19-2suse-SMP (root@hostname) (gcc ver-
sion 3.2) #1 SMP Wed Sep 18 13:08:12 UTC 2002
We are running under VM (31 bit mode)
This machine has an IEEE fpu
On node 0 totalpages: 131072
zone(0): 131072 pages.
zone(1): 0 pages.
zone(2): 0 pages.
Building zonelist for node : 0
Kernel command line: ramdisk_size=65536 root=/dev/ram1 ro cio_msg=no
Highest subchannel number detected (hex) : 0012
Calibrating delay loop... 788.07 BogoMIPS
Memory: 507608k/524288k available (1552k kernel code, 0k re-
served, 389k data, 56k init)
Dentry cache hash table entries: 65536 (order: 7, 524288 bytes)
Inode cache hash table entries: 32768 (order: 6, 262144 bytes)
Mount-cache hash table entries: 512 (order: 0, 4096 bytes)
Buffer-cache hash table entries: 32768 (order: 5, 131072 bytes)
Page-cache hash table entries: 131072 (order: 7, 524288 bytes)
debug: Initialization complete
POSIX conformance testing by UNIFIX
Detected 2 CPU's
Boot cpu address  1
cpu 0 phys_idx=1 vers=FF ident=111111 machine=2064 unused=0000
cpu 1 phys_idx=2 vers=FF ident=111222 machine=2064 unused=0000
migration_task 0 on cpu=0
migration_task 1 on cpu=1
init_mach : starting machine check handler
Linux NET4.0 for Linux 2.4
Based upon Swansea University Computer Society NET3.039
Initializing RT netlink socket
mach_handler : ready
mach_handler : waiting for wakeup
Starting kswapd
kinoded started
aio_setup: num_physpages = 32768
aio_setup: sizeof(struct page) = 44
pty: 256 Unix98 ptys configured
RAMDISK driver initialized: 16 RAM disks of 65536K size 1024 blocksize
loop: loaded (max 16 devices)
md: md driver 0.90.0 MAX_MD_DEVS=256, MD_SB_DISKS=27
md: Autodetecting RAID arrays.
md: autorun ...
md: ... autorun DONE.
debug: cio_msg: new level 6
```

```
debug: cio_trace: new level 6
debug: cio_crw: new level 6
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
IP: routing cache hash table of 2048 buckets, 32Kbytes
TCP: Hash tables configured (established 65536 bind 65536)
Linux IP multicast router 0.06 plus PIM-SM
NET4: Unix domain sockets 1.0/SMP for Linux NET4.0.
RAMDISK: Compressed image found at block 0
Freeing initrd memory: 7181k freed
VFS: Mounted root (ext2 filesystem).
done doing movetotmpfs
IPv6 v0.8 for NET4.0
IPv6 over IPv4 tunneling driver
Mounted /proc
Creating /var/log/boot.msg
Enabling system logging...
Sep 19 08:00:05 suse syslogd 1.4.1: restart.
```

Output 16: The Kernel Messages during the Boot Sequence

First, select the type of network connection to use: Token Ring, Ethernet, Ex-press Gigabit Ethernet, CTC, ESCON, IUCV, Hipersockets, or Express High-Speed Token Ring (Output 17). In this example installation, CTC is used, so 4 is chosen.

```
=                                                     =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                     =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA-Express High-Speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 17: Supported Network Connection Types

The system now displays a possible CTC configuration. If you prefer another setup, enter the device address of the CTC read channel (600 in our example) and then the one of the CTC write channel (601).

The protocol mode switches to one of several modes supported by CTC. By default, mode `0` (Compability mode) is selected.

> ┌─ **Note** ───
>
> To use IPv6 with the CTC device, protocol mode `1` (Extended mode) must be selected. IPv6 is not supported in protocol mode `0`.
>
> ── **Note** ─┘

SuSE Linux Enterprise Server now tries to load the network module by building a parameter line with the information provided then displays all loaded modules. Loading was successful if you get an output like the one in Output 18.

```
List of first 10 CTC Channels that were detected:
Dev. Channel type  CHPIDS
0600 CTC or LCS    0000000000000000
0601 CTC or LCS    0000000000000000
Device address for read channel (600):
Device address for write channel (0601):
Select protocol number for CTC:
0) Compatibility mode, also for non-Linux peers other
   than OS/390 and z/OS (this is the default mode)
1) Extended mode
3) Compatibility mode with OS/390 and z/OS
Enter your choice (0):
Writing 'ctc0,0x0600,0x0601,0,0' to /proc/chandev
Starting ctc driver:
Warning: loading /lib/modules/2.4.19-2suse-SMP/kernel/drivers/s390/net/fsm.o
will taint the kernel: no license
  See http://www.tux.org/lkml/#export-tainted for
information about tainted modules
Module fsm loaded, with warnings
CTC driver Version: 1.55.10.1 with CHANDEV support initialized
ctc0: read: ch 0600 (irq 000b), write: ch 0601 (irq 000c) proto: 0
Module ctc loaded, with warnings
ctc                 46140   0 (unused)
fsm                  1888   0 ctc
ctc0 detected.
ctc0 is available, continuing with network setup.
```

Output 18: Network Device Driver Parameters

Next, enter the full host name, your IP address, and the IP address of your peer (the 'other' side of the CTC connection). CTC connections are always point to point connections (Output 19 on the next page). To install via Token Ring, Ethernet, or Gigabit Ethernet, enter additional information, like the netmask, broadcast and default gateway, but no peer IP address.

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

Refer to Chapter *Finishing Installation* on page 84, Section *CTC and Escon with the ctc Module* on page 91 to get an overview of the recommended MTU size.

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.20
Please enter the IP address of your peer:
 192.168.0.1
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain:
 cosmos.com
Please enter the MTU (Maximum Transfer Unit),
leave blank for default: (1500): 1492
```

Output 19: Network Device Name

Now a summary is displayed. Confirm whether your input is correct (Output 20). Before the network is started, enter a password that will be valid only during the installation. After having IPLed the installed system, enter the "real" root password. See Output 21 on the next page.

```
Configuration for ctc0 will be:
Full host name          : earth.cosmos.com
IP address              : 192.168.0.20
Peer IP address         : 192.168.0.1
DNS IP address          : 192.168.0.2
DNS search domain       : cosmos.com
MTU size                : 1492
Is this correct (Yes/No) ? Yes
```

Output 20: Network Settings Summary

```
For security reasons you have to set an temporary installation
system password for the user "root".
You'll be asked for it only when you telnet into this installation
system to limit the access to it and it will be cleared as soon
as you shut down or reset the installation system

Please enter the temporary installation password: suse
```

Output 21: Entering the Password

With all basic parameters set up, the network will be started. Check the output of ifconfig, which should contain two entries: a loopback (lo) connection and one connection out of tr0, eth0, ctc0, escon0, iucv0 or hsi0 with correct settings (Output 22). To ensure that all IP addresses are reachable, a ping is sent to the local IP, the gateway, and the name server.

```
Temporary installation password set.
restarting syslogd:Temporary installation password set.
ifconfig ctc0 192.168.0.20 pointopoint 192.168.0.1 mtu 1492
/sbin/ifconfig ctc0 :
ctc0      Link encap:Serial Line IP
          inet addr:192.168.0.20  P-t-P:192.168.0.1
          Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP  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:100
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
Trying to ping my IP address:
PING 192.168.0.20 (192.168.0.20): 56 data bytes
64 bytes from 192.168.0.20: icmp_seq=0 ttl=255 time=0.212 ms
...
 --- 192.168.0.20 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.111/0.153/0.212 ms
Waiting 6 seconds for connection with remote side.
...
ctc0: connected with remote side
Trying to ping the IP address of the peer:
PING 192.168.0.1 (192.168.0.1): 56 data bytes
64 bytes from 192.168.0.1: icmp_seq=1 ttl=60 time=0.571 ms
...
 --- 192.168.0.1 ping statistics ---
3 packets transmitted, 2 packets received, 33% packet loss
round-trip min/avg/max = 0.571/2.463/4.355 ms
Peer seems to respond to our pings, continuing.
Trying to ping the IP address of the DNS Server:
```

```
PING 192.168.0.2 (192.168.0.2): 56 data bytes
64 bytes from 192.168.0.2: icmp_seq=0 ttl=254 time=1.492 ms
...
 --- 192.168.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 1.492/14.177/39.174 ms
Network Setup finished, running inetd...

You should be able to login via telnet now, for ssh wait a few
seconds, temporary host keys (only for installation) are being
generated now:
...
Generation of temporary installation host keys finished.
After installation, new, different SSH keys will be generated.
You should be able to login via telnet and ssh now.
...
```

Output 22: Example ifconfig

Where Next

Proceed with Chapter *Connecting to the SuSE Linux Installation System*.

Connecting to the SuSE Linux Installation System

After the network setup has been successfully completed, specify the source of the installation data. Chapter *Making the Installation Data Available* on page 19 describes the various sources of input. You must reference the source of input (installation data) selected when working through that section. Currently, the options nfs, smbfs, and ftp are supported.

Enter the option of your choice and specify the source directory on the installation server when prompted to do so. Finally, choose the connection method to the installation system from VNC, X, or ssh. Choose ond of these options and proceed with the appropriate initiation of the installation process.

Initiating the Installation for VNC

- 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 connect with vncviewer. Alternatively, an URL is given here that you can enter into your Java enabled-browser to connect to the installation system.

- Start a VNC client application on your client system. Either use vncviewer or the VNC Java client and a Java-enabled web browser.

- 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:

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

- After the connection has been established, start installing SuSE Linux Enterprise Server with YaST2.

Initiating the Installation for X Windows

- Make sure the server allows the client (the system that is installed) to connect to the server. Do this with the following command:

```
earth:~ #  xhost ⟨client IP address⟩
```

- Enter the IP address of the machine with the X server.

- Wait until YaST2 opens then start installation.

Initiating the Installation for SSH or Telnet

⌐ **Note** ───

Connecting from a Unix or Linux System
Run ssh or telnet from an xterm. Other terminal emulations do not properly interact with the text-based YaST2 interface.

─────────────────────────────────────── **Note** ⌐

To connect to the installation system using ssh or telnet, execute one of the following commands:

```
tux@ftpserver@earth:/home/tux >  telnet earth
```

or

```
tux@ftpserver@earth:/home/tux >  ssh earth
```

If your workstation runs on Microsoft Windows, use the ssh and telnet client and terminal emulator "putty", which is available on CD 1 in the directory `/dosutils/putty/`.

You will get a login prompt (Output 23). Enter root and log in with your password.

```
Trying 192.168.0.20
Connected to earth.cosmos.com
Escape character is '^]'.

   You are connected to earth.cosmos.com, Kernel 2.4-version
          It's 17:06 on Tuesday, 24 September 2002
```

```
    >>>>> ---------------------------------------- <<<<<
      >>> SuSE Linux Enterprise Server 8 for  zSeries    <<<
    >>>>> ---------------------------------------- <<<<<

     Logon as "root" with your temporary installation password
                    to start the installation.


earth login: root
Password:
```

Output 23: Connecting via telnet

Finally, start YaST2 with:

```
earth:~ #  yast2
```

As you will be using the text-based YaST2, consult Chapter *YaST2 in Text Mode (ncurses)* on page 119.

Where Next

Proceed with Chapter *Installation* on page 61.

Installation

YaST2 Takes Over

Now the actual installation of SuSE Linux Enterprise Server starts with the YaST2 installation program. All YaST2 screens have a common format. All entry fields, lists, and buttons on the YaST2 screens can be accessed with your mouse or keyboard.

If you have started YaST2 in the text-based mode only keyboard access to YaST2 is provided. Find further information about text-based YaST2 installation in Chapter *YaST2 in Text Mode (ncurses)* on page 119.

Selecting a Language

SuSE Linux Enterprise Server and YaST2 are adapted to use the language selected. Choose your language (Figure 12.1 on the following page). YaST2 selects time zone expected to suit your language settings. If your mouse does not work, navigate with the arrow keys to the desired language then press (Tab) repeatedly until 'Next' is selected. Then press (↵).
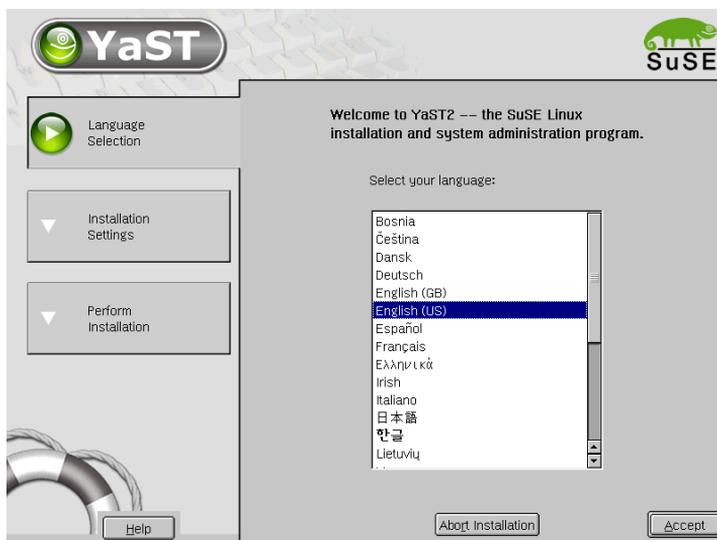
Figure 12.1: Selecting the Language

Making DASD Available for Installation

┌─ **Caution** ───

Low Level Formatting and Partitioning

YaST2 does not support unpartitioned DASD yet. Consult the README
located in the highest directory of SuSE Linux Enterprise Server CD 1
for directions if your DASD has not already been formatted and parti-
tioned.

── **Caution** ─┘

After a language is selected, a window appears, prompting for the parame-
ters to pass to the DASD driver (Figure 12.2 on the next page). A valid input
has the following syntax:

```
dasd=⟨DASD range⟩
```

For ⟨*DASD range*⟩, select the entire range of DASDs available or a few particular ones. With

```
dasd=fd00,fd03-fd08,fd46
```

make DASDs fd00, fd03, fd04, fd05, fd06, fd07, fd08 and fd46 available to the SuSE Linux Enterprise Server installation. Separate individual DASD device numbers with commas. Define ranges by placing a dash between the lowest and the highest address of the range.

Caution

Shared Devices in the IOCDS
Do not make shared devices available to SuSE Linux Enterprise Server that are already used by another operating system (e.g., OS/390 or SuSE Linux Enterprise Server). This may destroy any data on this device. Check your IOCDS before making shared devices available to SuSE Linux Enterprise Server.
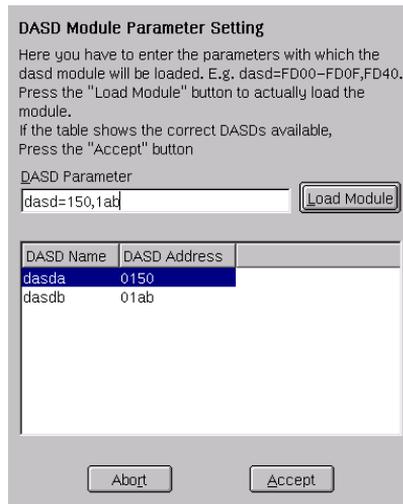

Caution

Figure 12.2: Making DASD Available for Installation

Press 'Load Module' to obtain a summary of the settings made. Leave this dialog by pressing 'Accept'.
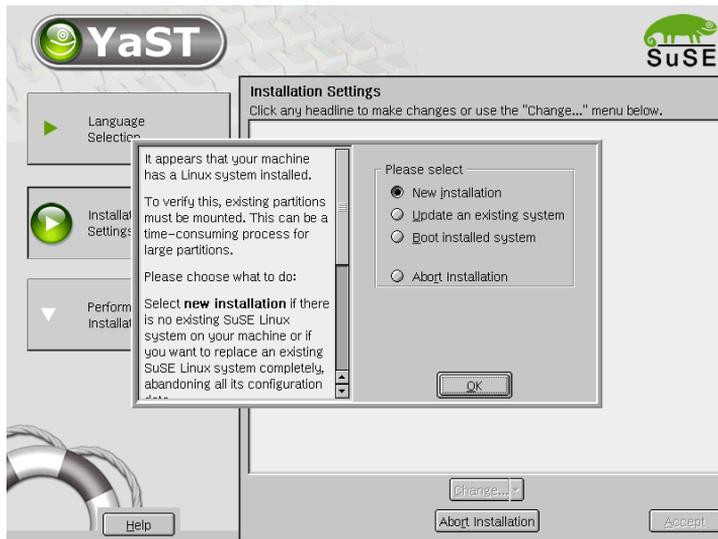
Installation Mode

Figure 12.3: Selecting the Installation Type

Choose 'New Installation' (Figure 12.3), which is the only valid option for this release of SuSE Linux Enterprise Server for S/390 and zSeries. 'Update an existing system' is not supported for SuSE Linux Enterprise Server 8.

Installation Suggestions

After the hardware has been detected, YaST2 provides information about the detected hardware and the suggestions for the installation and the partitioning (see Figure 12.4 on the next page). After you modify a suggestion, YaST2 returns to the suggestion window. The following sections explain the various configuration settings available.

Mode

Here, change the installation mode selected before the suggestion window appeared if you already have a Linux system on your computer.
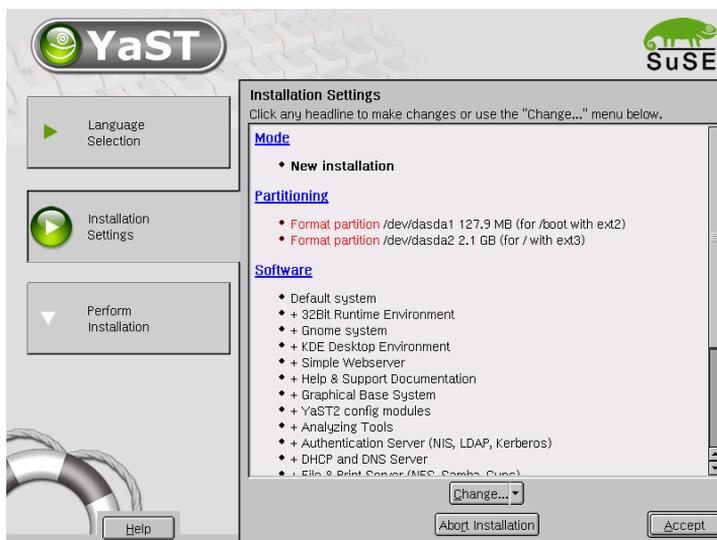
Figure 12.4: YaST2 Installation Settings

Partitioning

During installation, assign the disk space available to separate partitions. YaST2 assists in this task. SuSE Linux Enterprise Server for S/390 and zSeries discerns two different kinds of disks.

DASD A DASD can be formatted either in "CDL" (Compatible Disk Layout) or in "LDL" (Linux Disk Layout). CDL supports up to three partitions per DASD. LDL does not allow more than one partition per DASD.

SCSI SCSI disks can hold up to fifteen partitions. If your SCSI disk should be divided in more than four partitions, make one of these partitions an extended one. The extended partition can hold any desired number of logical partitions.

Manual Partitioning

With the 'Partitioner', shown in Figure 12.5 on the following page, the partitions of your hard disk can be modified manually. Partitions can be added, removed, or changed.
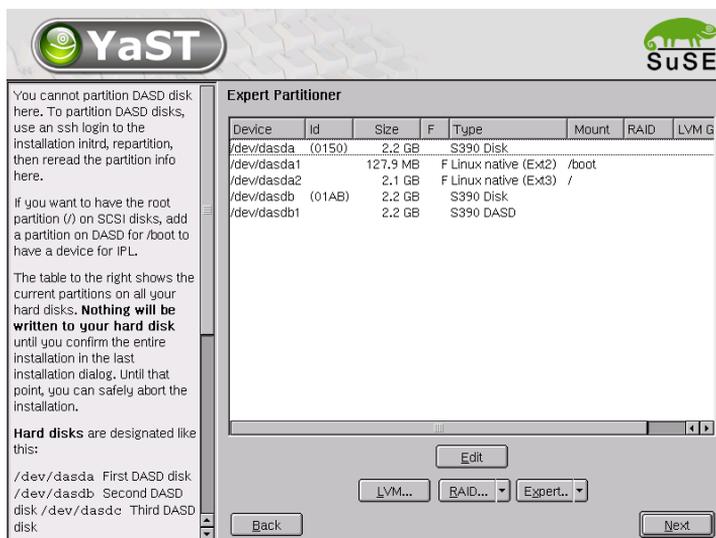
Figure 12.5: *The SuSE Linux Enterprise Server Partitioner*

If you select 'Partitioning' in the suggestion screen, the partitioner lists the hard disks and all available or suggested partitions. Disks are listed as devices without numbers (such as /dev/dasda or /dev/sda). Partitions are listed as parts of the devices (for example, /dev/dasda1 or /dev/sda1). The size, type, file system, and mount point are also displayed. The mount point describes where the partition is attached in the Linux file system tree.

Creating a Partition

To create a new partition:

1. Select the disk on which to create a partition (this will be done automatically if there is only one hard disk).

2. Select 'Create'. A dialog appears asking for the type of partition.

3. Select the file system to use to format the hard disk and, if required, a mount point. YaST2 suggests a mount point for each partition created. Details of the parameters are provided in the next section.

4. Select 'OK' to apply your changes.

The new partition is then listed in the partition table. Selecting 'Next' writes the partition table to disk and formats partitions if necessary.

Partitioning Parameters

To add a partition to the file system tree, set the following parameters in the partitioner:

1. Select the partition

2. 'Edit' the partition and set the parameters:

   - File system ID (if you want to format the partition): This can be 'Linux swap', 'Linux', 'Linux LVM', or 'Linux RAID'. For details on LVM and RAID, refer to Chapter *LVM — The Logical Volume Manager* on page 123 or to the *Administration* manual, respectively.

   - File system (for formatting the partition): This can be 'Swap', 'Ext2', 'Ext3', 'ReiserFS', or 'JFS'. Swap is a special format that makes the partition usable as virtual memory. Every system should have at least one swap partition of at least 128 MB. By default, ReiserFS is used for the partitions. Like JFS and Ext3, it is a "Journaling File System". A journaling file system enables a quick recovery from a system crash or a bad unmount. Also ReiserFS is very fast in handling lots of small files. Ext2 is a not journaling file system. It is "rock solid" and good for smaller partitions, as it does not require too much disk space for the management.

   - Mount Point: Set the directory where the new partition should be connected in the file system tree. If you select the last entry in the list, enter your own directory name. 'Swap' is used for the special file system Swap.

3. Select 'Next' to format and enable the partition.

Note ───

If you partition manually, create a swap partition. The swap partition is used to extend the available memory of your system.

── **Note** ─┘

More Partitioning Tips

If the partitioning is performed by YaST2 and other partitions are detected on the system, these partitions are entered in the file `/etc/fstab` to enable easy access to this data. This file contains all partitions in the system with their properties (parameters), such as the file system, mount point, and user permissions. An excerpt is shown in File 1.

```
/dev/sda1        /data1  auto     noauto,user 0 0
/dev/sda8        /data2  auto     noauto,user 0 0
```

File 1: /etc/fstab: Partition Data

The partitions are specified with the options `noauto` and `user`. This allows any user to mount or unmount these partitions if needed. For security reasons, YaST2 does not automatically enter the `exec` option here. However, to run programs from there, you can enter this option yourself. This measure will be necessary if you encounter system messages such as "bad interpreter" or "Permission denied".

Software

Here, determine which software to install on your machine.

'Minimal' This selection installs a fully functional Linux operating system providing text mode only.

'Default' This is the SuSE Linux Enterprise Server standard installation.

Choose 'Detailed selection' to do a more detailed software selection.

Preselection

The module starts with the selection filter. At the top left next to 'Filter', 'Selection' is marked. These selections represent groups of program packages to select for installation or removal by clicking the respective check box.
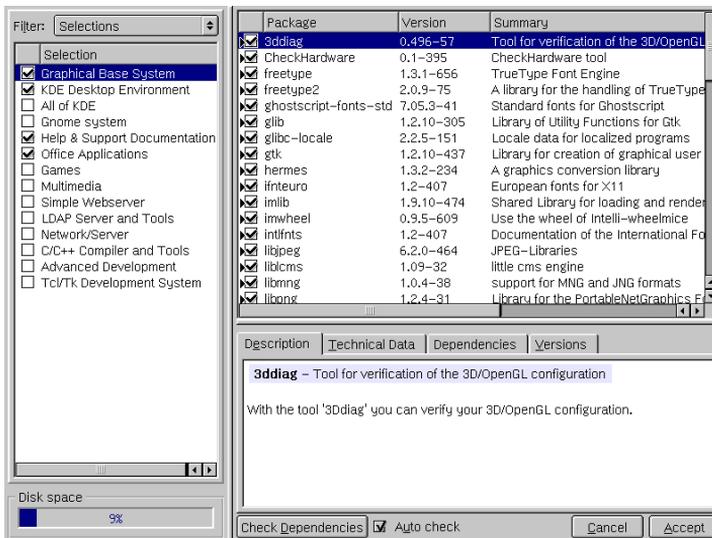
Figure 12.6: YaST2 Software Selection

Below, see possible preselection groups of this filter, some of which are already selected as they belong to the default installation of SuSE Linux Enterprise Server.

The right frame displays a list of the individual packages that belong to the selection. Packages selected for installation are checked to the left of the package name. Select and deselect individual packages according to your needs. To do this, click the symbol several times until the desired status is shown.

Other Filters

Click 'Filter' to see a selection of additional filters that can be used to structure the view of the packages. For example, there is a selection according to 'RPM groups', which is also defined as the default filter when you start the software selection in YaST2 after the system has been installed. Using this filter, the program packages are displayed according to subjects in a tree structure on the left side. The more you unfold the tree in a package group, the more detailed the selection will be and the smaller the number of related packages in the package list on the right side will be.

Time Settings

In this screen, shown in Figure 12.7, choose between `Local Time` and `GMT` in 'Set hardware clock to'. Your selection depends on the BIOS clock settings for your computer. If the hardware clock is set in the BIOS to `GMT`, SuSE Linux Enterprise Server automatically reflects Standard and Daylight Savings time changes.
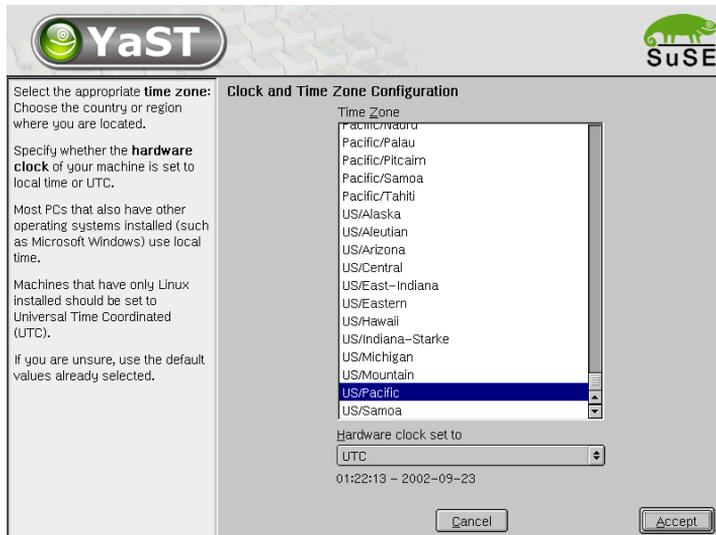

Figure 12.7: Selecting the Time Zone

Starting the Installation

Clicking 'Next' accepts the suggestions and any changes made. A confirmation screen opens. Click 'Yes' here to start the installation using your settings.

After the basic installation is finished, the system will shut down and needs to be IPLed from the boot device. The system will then proceed with the installation.

Where Next

For ESA Native or LPAR installation, proceed with Section *ESA Native Installation and LPAR Installation* on page 73 in Chapter *IPLing the Installed System*. For z/VM installation, proceed with Section *z/VM Installation* on page 73 in Chapter *IPLing the Installed System*.

IPLing the Installed System

ESA Native Installation and LPAR Installation

Now, IPL the installed system. Select 'LOAD' on the S/390 or zSeries terminal and select 'Clear'. Enter the load address (the device address of the root device) and start the load process.

z/VM Installation

Now, IPL the installed system. First, shut down the installation system with "halt". Log into the z/VM guest again as user "LINUX1" and IPL the installed system with

```
IPL 151 CLEAR
```

Where Next

Proceed with Chapter *Finishing the Installation* on page

Finishing the Installation

After the installed system has been IPLed, reconnect to the system to finish the installation. The actual steps to connect to the system depend on your initially selected connection method (see *Connecting to the SuSE Linux Installation System* on page 57).

Continuing the Installation for VNC

A message on the 3270 terminal asks you to connect with a VNC client to the Linux system. The message may be hidden between kernel messages and is easily missed or the terminal clears the screen before you can read it. In that case, connect the VNC viewer to the Linux system if nothing seems to happen for more than five minutes.

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

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

Continuing the Installation for the X Window System

Make sure the X11 server is running on the same system used to install the system when IPLing the installed system. A window with YaST2 opens and finishes the installation.

Continuing the Installation for SSH or Telnet

A message on the 3270 terminal asks you to connect with a VNC client to the
Linux system. The message may be hidden between kernel messages and is
easily missed or the terminal clears the screen before you can read it.

Perform the following steps to finish the installation:

- Use SSH to log in to the Linux system as user root. If you get a con-
  nection time-out or connection refused message, try again after a few
  minutes.

- Run the command

  ```
  linux@earth:  > /usr/lib/YaST2/bin/YaST2.sshinstall
  ```

  (not just "yast")

YaST2 will now be launched and the installation will continue with installing
the remaining packages (if any) and the initial system configuration.

System Configuration

After the installation of your system and selected software is complete, make
two more important settings before working with SuSE Linux Enterprise
Server: define a password for the system administrator root and create a
normal user. The following sections show how this is achieved.

root Password

The superuser (root) of a Linux system handles all kinds of administrative
tasks. root installs or removes software packages, changes the system config-
uration, and is responsible for hardware integration.

As a general rule, only log in as root to carry out administrative tasks, such as system maintenance or repairs. `root` is quite risky for everyday use, as `root` can delete files irreversibly.

For verification purposes, the password must be entered twice as shown in Figure 14.1. Be particularly careful not to forget the `root` password. It cannot be retrieved later.

> ## Caution
>
> **The User** `root`
> The user `root` has all the permissions needed to make changes to the system. To carry out such tasks, the `root` password is required. You cannot carry out any administrative tasks without this password.
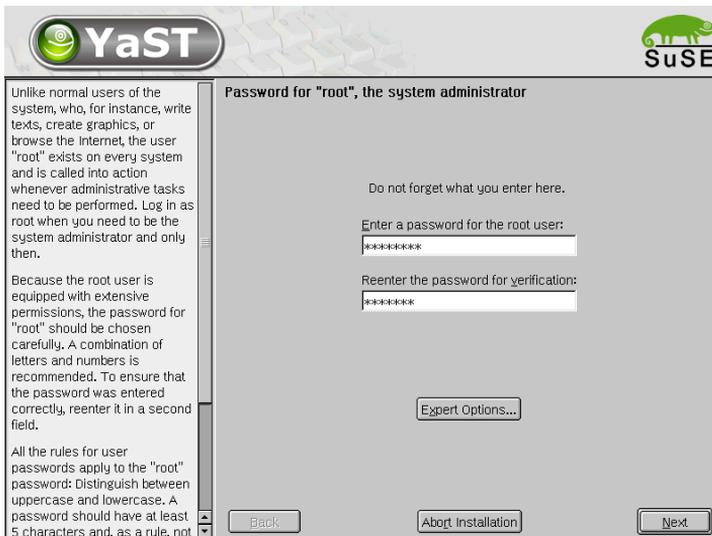>
> Caution

Figure 14.1: Setting the root Password

User Name and Password

Linux is an operating system that allows several users to work on the same system at the same time. Create a separate user account for any user who should be allowed to log in to your SuSE Linux Enterprise Server system.

Create such a user account yourself using the dialog as shown in Figure 14.2. Enter the user's first name and last name. Also specify the user name (login). If you cannot think of a suitable user name, click 'Suggestion' for the system automatically to generate one for you.
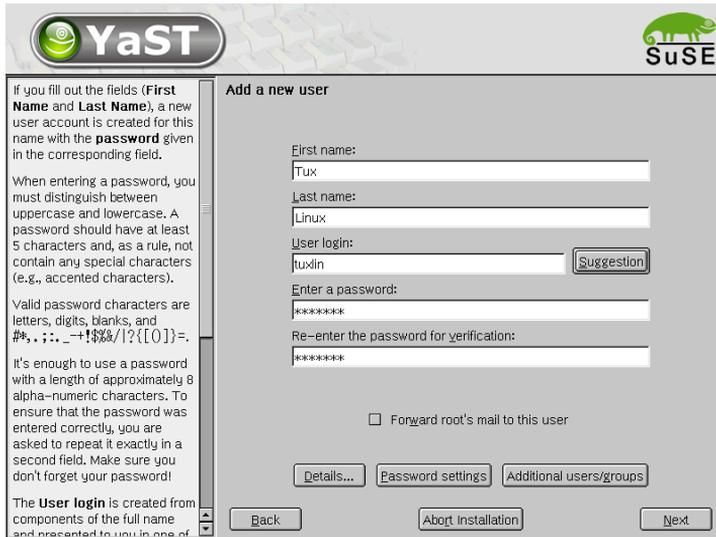
Figure 14.2: Entering the User Name and Password

Finally, enter a password for the user, which you must repeat to confirm. The user name tells the system who you are and the password verifies your identity.

Caution

Managing User Accounts

Memorize your user name and password, as you will need this information every time you log in. To provide effective protection, a password should be between five and eight characters long. The maximum length for a password is 128 characters. However, if no special modules are loaded, only the first eight characters are used to identify the password. Linux distinguishes between lowercase and uppercase letters in the password. Accented characters are not allowed. Special characters (such as *, ., # , and ;) and the digits 0–9 may be used.

Caution

Hardware Configuration

If desired, configure additional hardware at installation time if you have already planned your printer and network device setup in advance. These tasks can also be easily performed at a later time.

The 'Installation Settings' module shows which hardware is available and awaits configuration. Click the item of your choice and start configuration.
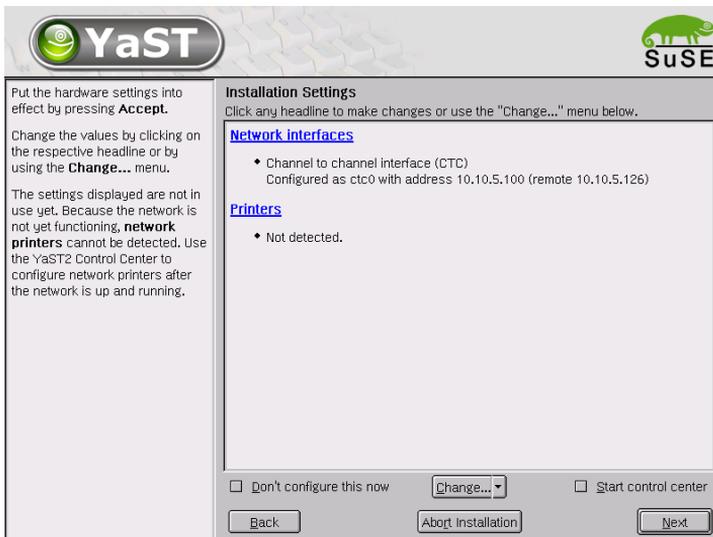

Figure 14.3: Configuring the System Components

Network Interfaces

If you have chosen the 'Network interfaces' option, a window appears showing an overview of both all configurable devices (listed in the upper part of the window) and the already configured ones (see bottom part of the window), as shown in Figure 14.4 on the following page. To configure another network device, select 'Other (not detected)' and click 'Configure'.

Once you have entered the 'Manual network card configuration' screen (Figure 14.5 on page 81), specify the network device. In SuSE Linux Enterprise Server S/390 and zSeries, choose from OSA Token Ring, OSA Ethernet, OSA-Express Gigabit Ethernet, Channel To Channel (CTC),
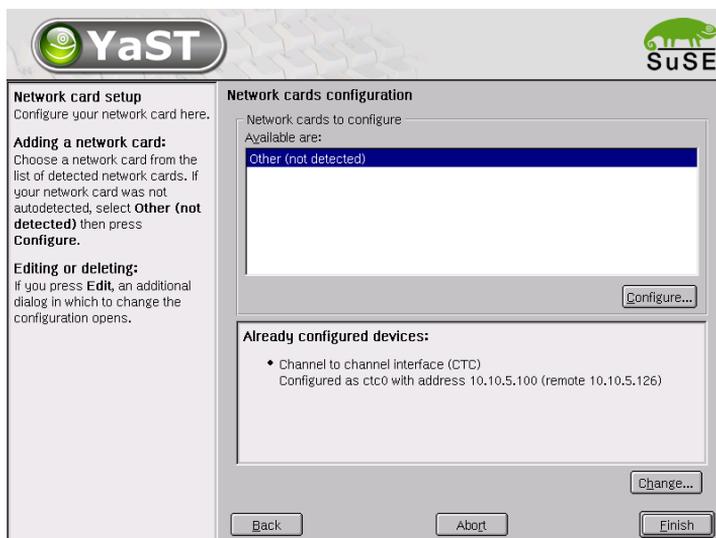
Figure 14.4: YaST2 Network Card Configuration

Escon, IUCV, Hipersockets, and OSA-Express High-Speed Token Ring. For the sake of simplicity, one example for an ethernet connection ("OSA Ethernet") and one example for a point-to-point connection ("CTC") are shown here.

> **Tip**
>
> Further information about network configuration on S/390 and zSeries can be retrieved from the *Device Drivers and Installation Commands* manual, which can be obtained at:
> http://www10.software.ibm.com/developerworks/
> opensource/linux390/june2003_documentation.shtml.
>
> **Tip**

Network Setup for OSA Ethernet

As you have specified the type of network device, YaST2 displays the kernel module to use. If necessary, specify some additional options for the network driver module here. Leave this mask with 'Next'. The next dialog prompts for the complete network address setup (Figure 14.6 on page 82).

For network address setup, first select the setup method for your network device (which is eth0 as you have chosen to set up an OSA Ethernet device).
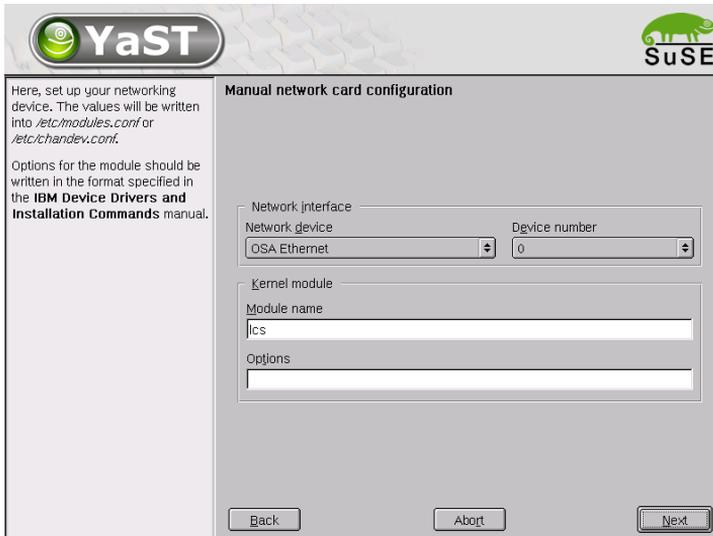
Figure 14.5: YaST2 Manual Network Card Configuration

Enter an IP address and the subnet mask for the network in which your machine will be located. Complete your network setup by configuring 'Routing' and 'Host name and name server'.

The 'Routing' dialog prompts for the default gateway to use for routing of all packages originating from your system. If you prefer to set up a more special and complex routing table, activate 'Expert configuration' and 'Add', 'Edit', or 'Delete' the desired route. Leaving this dialog by pressing 'Next' returns to the 'Network address setup' screen in which to continue with 'Host name and name server' configuration.

Enter the host name and the domain name of your system. Finally, enter a list of name servers and search domains to query. Leave with 'Next' until the final mask asks you to choose 'Finish' for the completion of your network setup.

Network Setup for CTC

Network setup for CTC ("Channel to Channel") closely resembles the setup procedure outlined for OSA Ethernet. Select the type of network device then accept or modify YaST2's choice of the kernel module and leave the 'Manual network card configuration' dialog using 'Next'. The 'Network address setup' dialog then opens, as shown in Figure 14.7 on page 83.
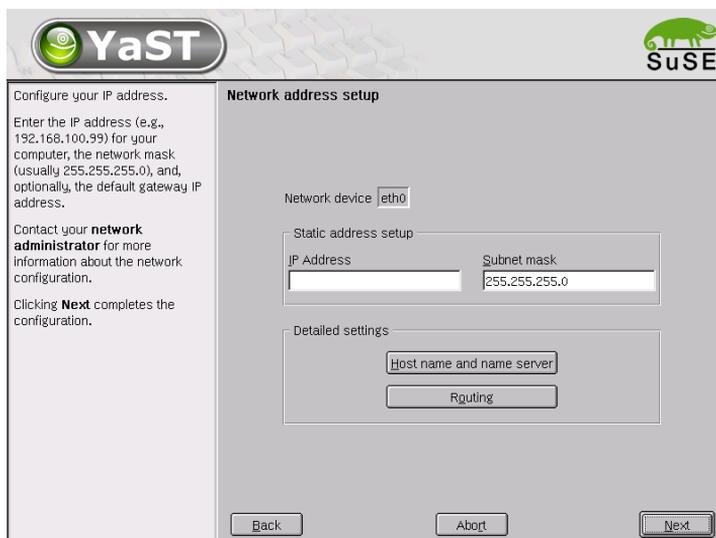
Figure 14.6: YaST2 Network Address Setup for OSA Ethernet

For static address setup, enter the IP address of your host and the IP address of its peer under 'Remote IP address'. Settings for routing, host name, and name server are made as described above.

Printer

For hardware reasons, SuSE Linux Enterprise Server for S/390 and zSeries does not support any kind of local printer connection. Print to files or choose to configure a remote printer queue over network.

Click 'Printer' in the 'Installation Setttings' screen to start autodetection of printers . To configure a network printer, select 'Others (not detected)' and press 'Configure'.

The following screen prompts for selection of the type of printer connection. See Figure 14.8 on page 84.

'IPP filtering and forwarding queue' is used as an example here. Confirm your choice by pressing 'Next' to enter the 'IPP Printer' dialog, which is shown in Figure 14.9 on page 85.
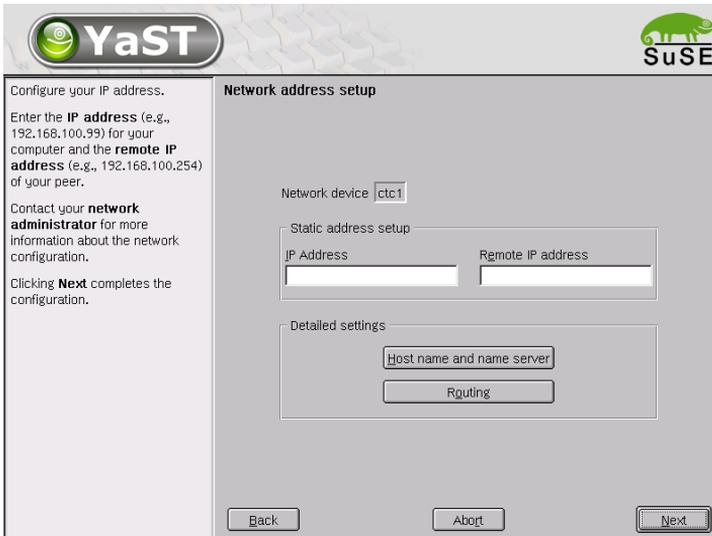
Figure 14.7: YaST2 Network Address Setup for CTC

Enter connection information, such as the host name of the printer server and the name of the remote printer queue. If you are not quite sure what to enter here, find these parameters from the network. To test the settings made, choose 'Test remote IPP access'.

Leave this dialog with 'Next', which continues to the 'Queue name' dialog. Enter the appropriate queue name and spooler settings. The next screen asks for specification of the manufacturer and model of the printer.

After entering the manufacturer and model of the printer, select a printer configuration. Leave this dialog with 'Next' and complete printer configuration by pressing 'Finish' in the autodetection screen. This returns to the 'Installation Settings' screen.

> **Tip**
>
> Find further in-depth information about printing under SuSE Linux Enterprise Server in Chapter *Printer Operation* in the *Administration* manual.
>
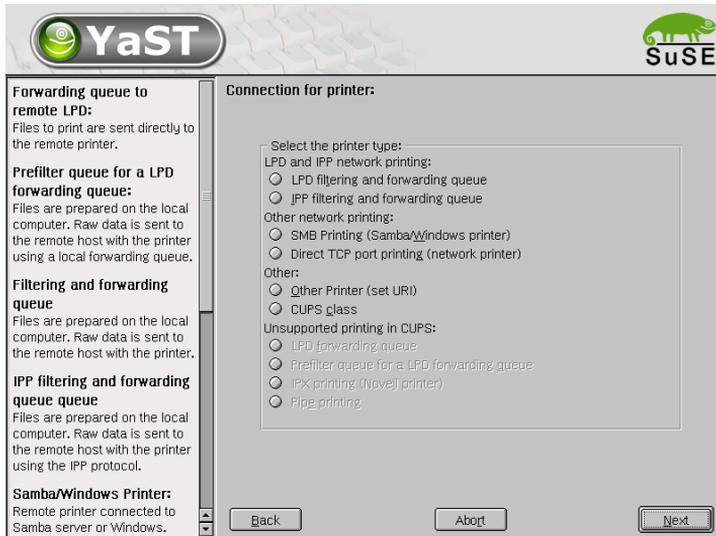> **Tip**

Figure 14.8: *Selecting the Printer Connection*

Finishing Installation

Click 'Next' if the system configuration is complete.

Tip

If you select 'Start control center', the YaST2 control center is started, allowing for further system configuration (i.e., configuration of NIS or YP)

Tip

The installation of SuSE Linux Enterprise Server for S/390 and zSeries is now complete. Start the YaST2 control center any time by simply logging in as root and executing the command `yast2`.
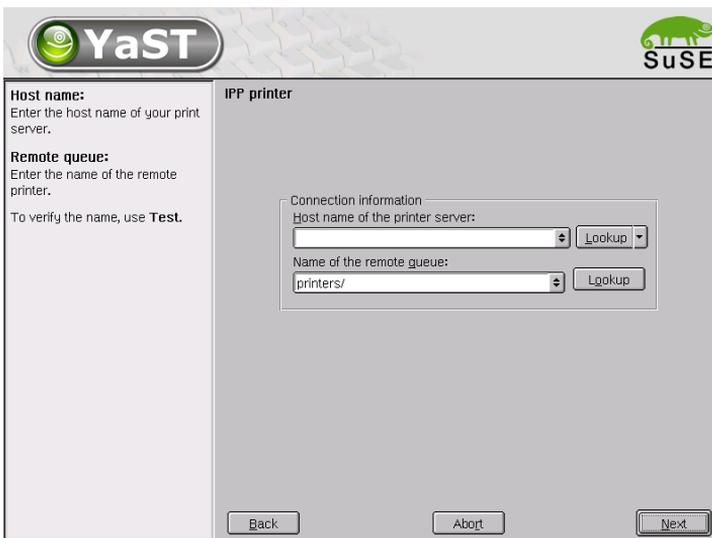
Figure 14.9: *Configuring IPP Printing*

Network Connection Types

SuSE Linux Enterprise Server Enterprise Server for S/390 and zSeries includes network drivers for OSA devices (Token Ring, Ethernet, and Gigabit Ethernet), CTC, Escon, Hipersockets, and IUCV. This chapter describes the configuration within the SuSE Linux Enterprise Server installation system.

Ethernet and Token Ring with the Ics Module

Select "1" from the list of network devices (Output 24) for a token ring network. For ethernet, select "2".

```
=                                                            =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                            =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA-Express High-speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 24: Supported Network Connection Types

When prompted, enter the network device number of the network card (e.g., FC20) and the relative port number on this device (Output 25).

```
Enter the read channel device number, e.g. 'FC20' (0c30):
fc20
Please enter the relative port number on device address fc20
Relative port, e.g. '0' (0):
2
Writing 'noauto;lcs0,0xfc20,0xfc21,0,1' to /proc/chandev
Using /lib/modules/2.4.version/net/lcs.o
Starting lcs module
with chandev support,with multicast support, with ethernet sup-
port,
with token ring support.
debug: lcs: new level 0
eth0: ip v6 supported yes enabled yes
eth0: multicast supported yes enabled yes
lcs: eth0 configured as follows read subchannel=162
write subchannel=163
read_devno=fc20 write_devno=fc21
hw_address=00:04:AC:E3:53:88 rel_adapter_no=1
lcs                    27584    0 (unused)
eth0      Link encap:Ethernet   HWaddr 00:04:AC:E3:53:88
BROADCAST MULTICAST   MTU:1500   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:100
RX bytes:0 (0.0 b)   TX bytes:0 (0.0 b)
eth0 is available, continuing with network setup.
```

Output 25: Network Device Driver Parameters

Next, enter the full host name, the IP address, and the netmask as well as the broadcast address, the IP address of the gateway, and the IP address of the DNS server. Finally, enter the DNS search domain and the MTU size (Output 26).

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.20
Please enter the net mask:
 255.255.255.0
Please enter the broadcast address:
 192.168.0.255
```

```
Please enter the gateway's IP address:
 192.168.0.1
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain:
 cosmos.com
Please enter the MTU (Maximum Transfer Unit),
leave blank for default: (1500):
```

Output 26: Network Device Name

A summary is displayed. Confirm whether your input is correct (Output 27).

```
Configuration for eth0 will be:
Full host name        : earth.cosmos.com
IP address            : 192.168.0.20
Net mask              : 255.255.255.0
Broadcast address     : 192.168.0.255
Gateway address       : 192.168.0.1
DNS IP address        : 192.168.0.2
DNS search domain     : cosmos.com
MTU size              : 1492
Is this correct (Yes/No) ? Yes
```

Output 27: Network Settings Summary

Gigabit Ethernet with the qeth Module

Select "3" from the list of network devices (Output 28).

```
=                                                     =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                     =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
```

```
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA High speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 28: Supported Network Connection Types

When prompted, enter the network device number of the network card (e.g.,
0x0808,0x0809, 0x080a) (Output 29).

Specify three device addresses (read, write, status) and the port name.

```
...
Possible configuration: read: 0x0808, write: 0x0809, data: 0x080a
Enter the device addresses for the qeth module, e.g. '0x0808,
0x0809,0x080a'
(0x0808,0x0809,0x080a): 0x0808,0x0809,0x080a
Please enter the portname (case sensitive) to use (suselin7):
SUSEPORT
Writing 'reset_conf' to /proc/chandev
Writing 'noauto;qeth0,0x0808,0x0809,0x080a;add_parms,0x10,0x0808,
0x080a,portname:SUSEPORT' to /proc/chandev
Using /lib/modules/2.4.version/net/qeth.o
qeth: loading qeth S/390 OSA-Express driver
qeth: Trying to use card with devnos 0x808/0x809/0x80A
qdio : CHSC returned cc 2. Using all SIGAs for sch xd.
qeth: Device 0x808/0x809/0x80A is an OSD Express card (level: 0223)
with link type Fast Eth (portname: FEF400)
qeth                 173040   0 (unused)
qdio                  47664   1 qeth
eth0 detected!
eth0      Link encap:Ethernet  HWaddr 00:02:55:09:D2:5D
          NOARP 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:100
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Interrupt:11

eth0 is available, continuing with network setup.
```

Output 29: Network Device Driver Parameters

Next, enter the full host name, the IP address, and the netmask as well as the
broadcast address, the IP address of the gateway, and the IP address of the
DNS server. Finally, enter the DNS search domain and the MTU (Output 30).

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.20
Please enter the net mask:
 255.255.255.0
Please enter the broadcast address:
 192.168.0.255
Please enter the gateway address:
 192.168.0.1
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain (e.g. suse.com):
 cosmos.com
Please enter the MTU (Maximum Transfer Unit,
leave blank for default) [1492]:
```

Output 30: *Network Device Name*

A summary is displayed. Confirm whether your input is correct (Output 31).

```
Configuration for eth0 will be:
Full host name          : earth.cosmos.com
IP address              : 192.168.0.20
Net mask                : 255.255.255.0
Broadcast address       : 192.168.0.255
Gateway address         : 192.168.0.1
DNS IP address          : 192.168.0.2
DNS search domain       : cosmos.com
MTU size                : 1492
Is this correct (Yes/No) ? Yes
```

Output 31: *Network Settings Summary*

CTC and Escon with the ctc Module

Select "4" from the list of network devices (Output 32) for CTC connection. For Escon, select "5". The driver is the same. The only difference to the user is the name of the device (ctc0 for CTC and escon0 for Escon).

```
=                                                      =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                      =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA High speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 32: Supported Network Connection Types

CTC and Escon connections always come with pairs of device addresses. Enter the device address of the read and write channel. Additionally, the protocol mode for the CTC device must be set. By default, protocol mode 0 (Compability mode) is selected.

> **Note**
>
> To use IPv6 with the CTC device, protocol mode 1 (Extended mode) must be selected. IPv6 is not supported in protocol mode 0.
>
> **Note**

```
List of first 10 CTC Channels that were detected:
Dev. Channel type  CHPIDS
0600 CTC or LCS    0000000000000000
0601 CTC or LCS    0000000000000000
Device address for read channel (600):
Device address for write channel (0601):
Select protocol number for CTC:
0) Compatibility mode, also for non-Linux peers other
   than OS/390 and z/OS (this is the default mode)
1) Extended mode
3) Compatibility mode with OS/390 and z/OS
Enter your choice (0):
```

```
Writing 'ctc0,0x0600,0x0601,0,0' to /proc/chandev
Starting ctc driver:
Warning: loading
/lib/modules/2.4.19-2suse-SMP/kernel/drivers/s390/net/fsm.o
will taint the kernel: no license
See http://www.tux.org/lkml/#export-tainted for information about
tainted modules
Module fsm loaded, with warnings
CTC driver Version: 1.55.10.1 with CHANDEV support initialized
ctc0: read: ch 0600 (irq 000b), write: ch 0601 (irq 000c) proto: 0
Module ctc loaded, with warnings
ctc                    46140   0 (unused)
fsm                     1888   0 ctc
ctc0 detected.
ctc0 is available, continuing with network setup.
```

Output 33: Network Device Driver Parameters

Next, enter the full host name, the IP address, the netmask, the IP address of the peer (the 'other' side of the connection), and the IP address of the DNS server. Finally, enter the DNS search domain and the MTU size (Output 34).

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.20
Please enter the IP address of your peer:
 192.168.0.1
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain:
 cosmos.com
Please enter the MTU (Maximum Transfer Unit),
leave blank for default: (1500): 1492
```

Output 34: Network Device Name

A summary is then displayed. Confirm whether your input is correct (Output 35).

```
Configuration for ctc0 will be:
Full host name   : earth.cosmos.com
IP address       : 192.168.0.20
Peer IP address  : 192.168.0.1
DNS IP address   : 192.168.0.2
DNS search domain: cosmos.com
MTU size         : 1492
Is this correct (Yes/No) ? yes
```

Output 35: Network Settings Summary

IUCV with the netiucv Module

Select "6" from the list of of network devices (Output 36) for the IUCV de-
vice driver. Just like CTC, IUCV provides virtual point-to-point connections.

```
=                                                       =
==- Welcome to SuSE Linux Enterprise Server 8 for zSeries -==
=                                                       =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA High speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 36: Supported Network Connection Types

Enter the name of the IUCV peer. Leave it blank to use the default "$TCPIP"
user (Output 37).

```
The name of the peer guest, e.g. 'TCPIP' (TCPIP): TCPIP
Trying to start the netiucv module now...
insmod netiucv iucv='TCPIP' :
Using /lib/modules/2.4.version/kernel/drivers/s390/net/netiucv.o
netiucv                21776   0 (unused)
netiucv module is loaded, checking interface iucv0.
iucv0     Link encap:Serial Line IP
          POINTOPOINT NOARP  MTU:9216  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:50
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

iucv0 seems to exist, continuing with network setup.
```

Output 37: Network Device Driver Parameters

Next, enter the full host name, the IP address, the netmask, the IP address of the peer, and the IP address of the DNS server. Finally, enter the DNS search domain and the MTU (Output 38).

```
Please enter your full host name:
 earth.cosmos.com
Please enter your IP address:
 192.168.0.18
Please enter the IP address of your peer:
 192.168.0.17
Please enter the IP address of the DNS server
or 'none' for none:
 192.168.0.2
Please enter the DNS search domain:
 cosmos.com
Please enter the MTU (Maximum Transfer Unit),
leave blank for default: (1500):
```

Output 38: Network Device Name

A summary is displayed. Confirm whether your input is correct (Output 39).

```
Configuration for iucv0 will be:
Full host name   : earth.cosmos.com
IP address       : 192.168.0.18
Peer IP address  : 192.168.0.17
DNS IP address   : 192.168.0.2
DNS search domain: cosmos.com
MTU size         : 1500
Is this correct (Yes/No) ? Yes
```

Output 39: Network Settings Summary

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 YaST2 running there. User interaction is thus limited to the actual YaST2 installation controlled by YaST2 dialogs.

The following parameters can be passed to the installation routine, which will take 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 ten. Specify more than one parameter on a line. Separate the parameters by spaces. You may specify the parameters in any order. Always keep the PARAMETER=value string together in one line. Example:

```
IP_HOST=s390zvm01.suse.de IP_ADDR=10.11.134.65
```

Some of the following parameters are required. If they are missing, the automatic process pauses and ask for the value interactively.

General Parameters

- AUTOINSTALL=⟨*yes/no*⟩

  This parameter controls whether the other parameters are only default values that still must be acknowledged by the user. Set this parameter to *yes* if all values should be accepted and no questions asked.

- `INST_PASSWORD=⟨`*`your password`*`⟩`

  The temporary root password for access during the installation. The password must have between three and eight characters.

Configuring the Network Interface

- `IP_ADDR=192.168.55.23`

  Enter the IP address of the interface to configure.

- `IP_HOST=earth.cosmos.com`

  Enter the FQDN (Fully Qualified Domain Name) of the host to install.

- `IP_GATEWAY=192.168.55.254`

  Specify the gateway being used.

- `IP_DNS=192.168.123.123`

  Specify the DNS server in charge.

- `IP_SEARCH=example.com`

  Specify the search domain for DNS requests.

- `IP_MTU=1500`

  Specify an appropriate MTU size.

- `IP_INTERFACE=ctc`

  Enter the type of interface to configure. Possible values are `ctc`, `escon`, `iucv`, `eth`, `qeth`, `tr`, `hsi`, and `qeth-tr`.

  - ▷ For the interfaces of type `hsi`, `eth`, `qeth`, `qeth-tr`, and `tr`, specify an appropriate netmask and an optional broadcast address:
    `IP_NETMASK=255.255.255.0`
    `IP_BROADCAST=192.168.255.255`

▷ For the interfaces of type `ctc`, `escon`, and `iucv`, enter the IP address of the peer:
`IP_PEER=192.168.55.20`

- Each of the interfaces requires certain setup options:

  ▷ **Interface `ctc` and `escon`**

  `READ_DEVNO=0424`
  `WRITE_DEVNO=0425`

  ⟨*READ_DEVNO*⟩ specifies the READ channel to use.
  ⟨*WRITE_DEVNO*⟩ is the WRITE channel.

  ▷ For the `ctc` interface, specify the protocol that should be used for this interface:
  `PROTOCOL=`⟨*0/1/3*⟩

  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
3	Compatibility mode with OS/390 and z/OS

  ▷ **Interface `eth` and `tr`**

  `DEVNO=0124`
  `PORTNO=1`

  ⟨*DEVNO*⟩ stands for the channel number used in this setup. A second port number can be derived from this by adding one to ⟨*DEVNO*⟩. ⟨*PORTNO*⟩ is used to specify the relative port.

  ▷ **Interface `iucv`**

  `PEER=PARTNER`

  Enter the name of the peer machine.

  ▷ **Interface `qeth` for OSA-Express Gigabit Ethernet and `qeth-tr` for OSA-Express High speed Token Ring**

  `READ_DEVNO=0524`
  `WRITE_DEVNO=0525`
  `DATA_DEVNO=0526`
  `PORTNAME=FEF400`

  For ⟨*READ_DEVNO*⟩, enter the number of the READ channel.
  ⟨*DATA_DEVNO*⟩ specifies the DATA channel. For ⟨*PORTNAME*⟩,

enter an appropriate port name. Make sure the READ channel carries an even device number. The WRITE channel is derived from READ channel by adding one, so is not explicitly specified in this setup.

▷ **Interface** hsi for Hipersockets and VM guest LANs

```
READ_DEVNO=0624
WRITE_DEVNO=0625
DATA_DEVNO=0626
```

For ⟨*READ_DEVNO*⟩, enter the appropriate number for the READ channel. For ⟨*WRITE_DEVNO*⟩ and ⟨*DATA_DEVNO*⟩, enter the appropriate channel numbers.

Configuring DASDs

■ `DASD_Parameter=dasd=0150,01ab`

This parameter is passed to the DASD driver, which will then make the DASD available to the SuSE Linux Enterprise Server installation.

Specifying Source of Installation and YaST2 Interface

■ `INST_INFO=nfs`

Specify the type of installation to perform. Possible values are `nfs`, `SAMBA`, and `ftp`.

▷ If `INST_INFO=ftp` is chosen, specify the user and password for FTP installation. If these are not set, anonymous login is assumed:
`FTPUSER=`⟨*a user name*⟩
`FTPPASSWORD=`⟨*a password*⟩

▷ If `INST_INFO=SAMBA` is chosen, specify the user, password, and domain that should be used for SAMBA installation. If these three parameters are not set, guest login is assumed:
`SAMBAUSER=`⟨*a user name*⟩
`SAMBAPASSWORD=`⟨*a password*⟩
`WORKDOMAIN=`⟨*a domain*⟩

■ `INST_IP_ADDR=129.168.44.33`

Enter the IP address of the machine used as the installation server.

- INST_IP_DIR=/s390

  Enter a directory on the installation server, which holds the installation data of the SuSE Linux Enterprise Server CDs.

- INST_SCREEN=⟨*VNC/X/SSH*⟩

  Decide whether an X server should be used for installation or whether to use VNC or SSH installation. Enter X to use an X server and VNC or SSH for VNC or SSH installation.

  ▷ For X, specify the IP address of the X server to use for installation:
  INST_X_ADDR=10.10.10.10
  Enter the IP address of the X server. To enable a connection between YaST2 and the X server, run xhost ⟨*IP_ADDR*⟩ on this machine.

  ▷ For VNC, specify a password of six to eight characters to use for installation:
  VNC_PASSWORD=⟨*a password*⟩

Support for Fibre Channel–Attached SCSI

IBM zSeries can access Fibre Channel–attached SCSI devices. To exploit this feature, the zfcp (SCSI-over-Fibre Channel Protocol) device driver module must be loaded and configured. Using the zfcp device driver, SCSI devices on zSeries are accessible to Linux like any other SCSI device under Linux.

Loading and Configuring the zfcp Module Manually

> **Tip**
>
> Before creating your own zfcp configuration, refer to the IBM *Device Drivers and Installation Commands* Manual at http://oss.software.ibm.com/developerworks/opensource/linux390/docu/lx24jun03dd01.pdf.
>
> **Tip**

The zfcp configuration is located in /etc/zfcp.conf. Open this configuration file with vi /etc/zfcp.conf. Uncomment the last line and modify it according to your local configuration by specifying ⟨*devno*⟩, ⟨*Linux SCSI Port*⟩, ⟨*fcp wwpn*⟩, ⟨*Linux SCSI Lun*⟩, and ⟨*fcp-lun*⟩. A standard zfcp.conf is presented in Output 40.

```
# This file contains the zfc configuration as described
# in the section 'SCSI-over-Fibre Channel driver' of
# the Device Drivers and Installation Commands Manual.
#
# WARNING: In case you modify the configuration in this
# file after finishing the installation of SuSE Linux
# Enterprise Server, you have to save it to the boot
# process (i.e. initrd) before the next reboot to make
# it persistent. Failing in doing so may leave your
# system in a misconfigured state after the next reboot
# which - in the worst case - may lead to data loss.
#
# To store the zfcp configuration to the boot process,
# follow these steps:
#
# * prepare an initial ramdisk (initrd) containing
#   the configuration:
#                   mk_initrd
#
# * store the initrd to the boot loader:
#                   zipl
#
# This file will then be written to
# /proc/scsi/zfcp/add_map at system boot.
#
#   <devno> <Linux SCSI Port>:<fcp wwpn> <Linux SCSI Lun>:<fcp-lun>
#   devno = fcp adapter device number
#   fcp = fibre channel protocol
#   wwpn = world wide port number
```

```
#    lun = logical unit number
#
#    always modify and uncomment this to match your local configuration:
#    0x1234 0x1:0x345adf3322443525 0x0:0x01f000000000000
```

Output 40: An Example /etc/zfcp.conf

After the configuration has been modified and saved, load the kernel modules:

```
earth:~ #  modprobe scsi_mod
```

```
earth:~ #  modprobe zfcp
```

Configure the kernel with your configuration:

```
earth:~ #  cat /etc/zfcp.conf > /proc/scsi/zfcp/add_map
```

Verify that your changes have been applied using the following command, which shows all configured zfcp-attached SCSI device mappings:

```
earth:~ #  cat /proc/scsi/zfcp/map
```

The SCSI modules for needed devices are loaded automatically when you access the corresponding SCSI device nodes (/dev/sd* for disks, /dev/sr* for CD-ROM, /dev/st* for tapes). You can also query the SCSI system about its state with cat /proc/scsi/scsi. If the attached devices fail to show up in the displayed list, you can force a SCSI-bus rescan with the command

```
earth:~ #  echo "scsi scan-new-devices" > /proc/scsi/scsi
```

Your SCSI device is now ready to use.

For SuSE Linux Enterprise Server to initialize the zfcp driver according to your configuration at boot, make your settings persistent by saving your /etc/zfcp.conf to the boot process. This is done with the following commands:

```
earth:~ #  mkinitrd
```

```
earth:~ #  zipl
```

Partitioning and formatting Fibre Channel–Attached SCSI Disks

You need to create a partitioning on the SCSI disk prior to use with SuSE Linux Enterprise Server. This is done with the command fdisk. Assuming that your SCSI disk can be accessed via /dev/sda you can create a partition with the command

```
earth:~ #  fdisk /dev/sda
```

As an example we will create a new primary partition on an otherwise unpartitioned SCSI disk. To create a new primary partition you need to press 'n' (for new partition), 'p' (for primary partition), and then specify the first cylinder and the last cylinder of the new partition. 'p' (for print) prints the current partition table, and 'w' writes the partition table to the disk.

Following these steps the output on your display will look like this:

```
Command (m for help): p

Disk /dev/sda: 62 heads, 62 sectors, 1016 cylinders
Units = cylinders of 3844 * 512 bytes

   Device Boot     Start        End     Blocks   Id  System

Command (m for help): n
Command action
   e    extended
   p    primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-1016, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1016, default 1016):
Using default value 1016

Command (m for help): p

Disk /dev/sda: 62 heads, 62 sectors, 1016 cylinders
Units = cylinders of 3844 * 512 bytes

   Device Boot     Start        End     Blocks   Id  System
/dev/sda1             1        1016    1952721   83  Linux

Command (m for help): w
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
Syncing disks.
```

Output 41: *Output from fdisk*

Now you need to create a filesystem on the newly created partition. Assuming that your partition is on device /dev/sda1 and you want to create an ext3–filesystem you can use the command

```
earth:tld #   mke2fs -j /dev/sda1
```

The following output will be printed:

```
mke2fs 1.28 (31-Aug-2002)
Filesystem label=
OS type: Linux
Block size=4096 (log=3D2)
Fragment size=4096 (log=3D2)
244320 inodes, 488180 blocks
24409 blocks (5.00%) reserved for the super user
First data block=0
15 block groups
32768 blocks per group, 32768 fragments per group
16288 inodes per group
Superblock backups stored on blocks:
        32768, 98304, 163840, 229376, 294912

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 28 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

Output 42: *Output from mke2fs*

To use another filesystem (e.g. reiserfs) you need to use the appropriate command (e.g. mkreiserfs) instead of mke2fs.

Installing on zfcp-Attached SCSI Disks

Direct installion on fibre channel–attached SCSI Disks is currently not support. To install on a SCSI device, you should first install SuSE Linux Enterprise Server on a DASD-Device as described previously and then copy the installed programs and files to the SCSI device.

First you need to prepare the SCSI device as described in the previous chapter. This device need to be mounted to be accessible for installation. Assuming the device should be mounted under the mount–point /mnt you need to issue the command

```
earth:tld #   mount /dev/sda1 /mnt
```

Now you need to create the same file structure as on your original installation device. For this you first need to create some directories:

```
earth:tld #   cd /mnt
earth:tld #   mkdir mnt proc tmp
```

and then copy all files and directories to the new disk:

```
earth:tld #   cp -r -no-dereference -preserve=all -target-
    directory=/mnt /bin /dev /home /lib /lib64 /sbin /var /boot
    /etc /root /srv /usr /opt
```

The directory structure on the SCSI device should now look like this:

```
.    bin   dev  home  lib64       mnt  proc  sbin  tmp  var
..   boot  etc  lib   lost+found  opt  root  srv   usr
```

Output 43: Directory layout of the SCSI device

Finally we need to change the permissions for the /tmp directory to be readable and writeable for all, as everybody should be able to use this directory:

```
earth:tld #   chmod 1777 /mnt/tmp
```

Your SuSE Linux Enterprise Server system is now installed on the SCSI device. As a last step you need to configure the file system table /mnt/etc/fstab to correctly reference your SCSI device.

Assuming that `/mnt/etc/fstab` looks like this:

```
/dev/dasda2      /            reiserfs   defaults             1 1
/dev/dasda1      swap         swap       pri=42               0 0
devpts           /dev/pts     devpts     mode=0620,gid=5      0 0
proc             /proc        proc       defaults             0 0
```

Output 44: Original /etc/fstab

you need to change the first line to reference the SCSI device as a root disk:

```
/dev/sda2        /            ext3       defaults             1 1
/dev/dasda1      swap         swap       pri=42               0 0
devpts           /dev/pts     devpts     mode=0620,gid=5      0 0
proc             /proc        proc       defaults             0 0
```

Output 45: Modified /etc/fstab

IPLing from an zfcp-Attached SCSI Disk

Caution

IPLing from a zfcp-attached SCSI Disk

IPLing from zfcp-attached disks is a new feature of z/VM 4.4.0. Earlier releases will *not* work as the necessary functions are missing. You should also make sure that the MicroCode Levels required to IPL from zfcp are installed (MCLs not yet available)

Caution

After your system has been installed on the SCSI disk as described in the previous chapter we need to check the configuration data for IPLing from zfcp-devices and write the IPL record to the SCSI disk.

First you need to check the parameter file in `/mnt/boot/zipl/parmfile` whether it refers to the SCSI disk. It should look like this (The DASD part is not really necessary).

```
dasd=150,1ab
root=/dev/sda1 ro
```

Output 46: An Example /boot/zipl/parmfile

where /dev/sda1 should be replaced with the correct device name in case your system is installed on a different SCSI disk partition.

Next, you need to adapt the "zipl" configuration file /mnt/etc/zipl.conf to refer to your SCSI disk. It should look like this:

```
# Generated by YaST2
[defaultboot]
default=ipl

[ipl]
target=/boot/zipl
image=/boot/kernel/image
ramdisk=/boot/initrd
parameters="root=/dev/sda1"
```

Output 47: Modified /mnt/etc/zipl.conf

To make sure that you only change the boot environment for the system installed on the SCSI disk it is recommended to temporarily switch the root directory to the SCSI disk. This is done with the command "chroot"

```
earth:tld #  chroot /mnt
```

Now you should rebuild /boot/initrd and re-write the IPL record with the commands:

```
earth:~ #  mkinitrd
earth:~ #  zipl
```

Finally you can return to your original root device with

```
earth:tld #  exit
```

The system is now ready for IPLing from the SCSI device.

ESA Native Installation and LPAR Installation

Now you can IPL the installed system. Select 'LOAD' on the S/390 or zSeries SE (Service Element) panel and select 'load type' SCSI.

Now you have to fill in the parameter from your /etc/zfcp.conf.

Load address: the zfcp device number

World wide port name: the zfcp wwpn

Logical unit number: the zfcp lun

Boot program selector: If you have more than one entry in the 'menue:' section of your /etc/zipl.conf, you can enter the number of the menue entry here. Valid numbers are 0 - 30; 0 is the default.

Click on 'OK' to start IPL.

z/VM Installation

To start IPL under z/VM you first have to set the load-device to be used for IPL with the command:

```
SET LOADDEV PORT <zfcp-wwpn> LUN <zfcp-lun> BOOT <num>
```

where you have to replace the appropriate values for ⟨*zfcp-wwpn*⟩, ⟨*zfcp-lun*⟩. Both of these values should be given in hexadecimal; at least one separator blank is needed after the 8th character if the number is longer than 8 hex characters.

⟨*num*⟩ is the number of the corresponding menue entry in /etc/zipl.conf, if another than the default 1st entry is to be used.

You can query the existing value for LOADDEV with the command

```
QUERY LOADDEV
```

You can now start IPL with

```
IPL <fcp-adapter-devno>
```

Using initrd as a Rescue System

If the kernel of the SuSE Linux Enterprise Server for S/390 and zSeries is upgraded or modified, it is possible to reboot the system accidentally in an inconsistent state, so standard procedures of IPLing the installed system fail. This most commonly occurs if a new or updated SuSE Linux Enterprise Server kernel has been installed and the "zipl" program has not been run to update the IPL record. This chapter describes a method of using the standard installation package as a rescue system from which the "zipl" program can be executed to update the IPL record.

> **Note**
>
> For this method to work, the SuSE Linux Enterprise Server for S/390 and zSeries installation data must be available. For details, refer to Chapter *Making the Installation Data Available* on page 19. Additionally, you need the channel number of the device and the partition number within the device that contains the root file system of the SuSE Linux Enterprise Server installation.
>
> **Note**

IPLing the Rescue System

First, IPL the SuSE Linux Enterprise Server for S/390 and zSeries installation system as described in Chapter *IPLing the SuSE Linux Installation System* on page 39. A list of choices for the network adapter to use is then presented.

```
=                                                       =
==-  Welcome to SuSE Linux Enterprise Server 8 for zSeries   -==
```

```
=                                                               =

Please select the type of your network device:
 0) no network
 1) OSA Token Ring
 2) OSA Ethernet
 3) OSA-Gigabit Ethernet or OSA-Express Fast Ethernet
 4) Channel To Channel
 5) Escon
 6) IUCV
 8) Hipersockets
 9) OSA High speed Token Ring (qdio)
10) Show subchannels and detected devices
Enter your choice (0-10):
```

Output 48: *Network Type Selection*

Select 0 for *no network*. The installation program is terminated and the following messages are printed:

```
*** OK, NETWORK ACCESS WILL _NOT_ BE AVAILABLE. ***


You should be able to login via telnet/ssh now.
To restart network setup, enter:

netsetup

To continue the installation, enter:

./inst_source

bash: no job control in this shell
SuSE Instsys suse:/ #
```

Output 49: *Output for No Network Adapter Selected*

This opens a root shell from which to issue all necessary commands directly.

Loading DASD Modules

To access the root device, the required kernel modules must be loaded. First, the DASD modules must be loaded. They consist of a middle layer module "dasd_mod" and a low-level module depending on the DASD type. Here, ECKD devices are assumed, so the module is called "dasd_eckd_mod". The "dasd_mod" module requires the channel ranges of the DASDs to access (e.g., `0150`) as an argument (See Chapter *Making DASD Available for Installation* on page 62). The modules are loaded with the commands:

```
SuSE Instsys suse:/ # insmod dasd_mod dasd=0150
Using /lib/modules/version/kernel/drivers/s390/block/dasd_mod.o
dasd: initializing...
debug: dasd: new level 3
dasd: Registered successfully to major no 94
dasd: initialization finished
SuSE Instsys suse:/ # insmod dasd_eckd_mod
Using /lib/modules/version/kernel/drivers/s390/block/dasd_eckd_mod.o
dasd(eckd): ECKD discipline initializing
[ ... ]
Partition check:
 dasda:VOL1/  0X0150: dasda1 dasda2
dasd(eckd): We are interested in: CU 3880/00
dasd(eckd): We are interested in: CU 3990/00
dasd(eckd): We are interested in: CU 2105/00
dasd(eckd): We are interested in: CU 9343/00
```

Output 50: Loading DASD Modules

If the line `Partition check` is printed, all modules have been loaded properly and the DASD device is now available for mounting.

Mounting the Root Device

If the modules have loaded correctly, you should now be able to mount the root device. Assuming that the root device is on the second partition of the DASD device (`/dev/dasda2`) the corresponding command is

```
SuSE Instsys suse:/ # mount /dev/dasda2 /mnt
```

Output 51: Mounting the Root Device

> **Note**
>
> If the installed system has not been shutdown properly, it may be
> advisable to check the file system consistency prior to mounting.
> This prevents any accidental loss of data. In our example, issue the
> command `fsck /dev/dasda2` to ensure that the file system is in a
> consistent state.
>
> **Note**

By just issuing the command `mount`, it is possible to check whether the file
system could be mounted correctly.

```
SuSE Instsys suse:/ # mount
shmfs on /newroot type shm (rw,nr_inodes=10240)
devpts on /dev/pts type devpts (rw)
virtual-proc-filesystem on /proc type proc (rw)
/dev/dasda2 on /mnt type reiserfs (rw)
```

Output 52: Output of the Mount Command

Changing to the Mounted File System

For the "zipl" command to read the configuration file from the root device of
the installed system and not from the rescue system, change the root device
to the installed system. This is done with the `chroot` command:

```
SuSE Instsys suse:/ # cd /mnt
SuSE Instsys suse:/mnt # chroot /mnt
```

Output 53: chroot to the Mounted File System

Executing "zipl"

Now execute the command `zipl` to rewrite the IPL record with the correct
values:

```
sh-2.05b# zipl
building bootmap     : /boot/zipl/bootmap
adding Kernel Image : /boot/kernel/image located at 0x00010000
adding Ramdisk       : /boot/initrd located at 0x00800000
adding Parmline      : /boot/zipl/parmfile located at 0x00001000
Bootloader for ECKD type devices with z/OS compatible lay-
out installed.
Syncing disks....
...done
```

Output 54: Installing the IPL Record with zipl

Exiting the Rescue System

To exit the rescue system, first leave the shell opened by the chroot command:

```
sh-2.05b# exit
exit
SuSE Instsys suse:/mnt #
```

Output 55: Exit from the chrooted Shell

To prevent any loss of data, flush all unwritten buffers to disk with the sync command. Now change to the root directory of the rescue system and unmount the root device of the SuSE Linux Enterprise Server for S/390 and zSeries installation.

```
SuSE Instsys suse:/mnt # cd /
SuSE Instsys suse:/ # umount /mnt
```

Output 56: Unmounting the File System

Finally, halt the rescue system with the halt command. The SuSE Linux system can now be IPLed as described in Chapter *IPLing the Installed System* on page *73*.

YaST2 in Text Mode (ncurses)

YaST2 can also be used from a text-based terminal. This is useful when the administrator cannot access the system via a graphical console running X11.

Invocation and Usage

To start YaST2 in text mode, enter `yast` as `root` in a terminal.

The usage may be unfamiliar, but is very simple. The whole program can be operated with (Tab), (Alt) + (Tab), (Space), the arrow keys ((↑) and (↓)), and (Enter) as well as with shortcuts. When YaST2 is started in text mode, the YaST2 Control Center appears first, as shown in Figure D.1.
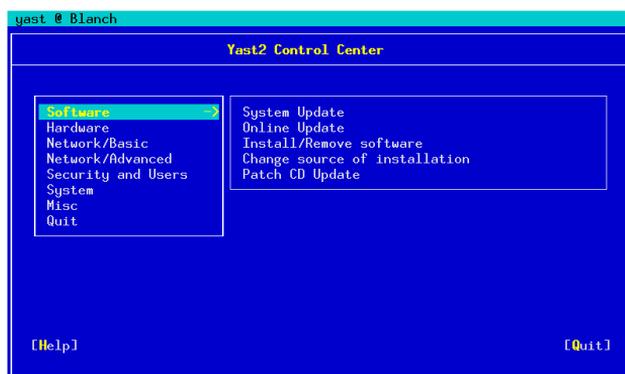
Figure D.1: Main Window of the YaST2 Control Center

The window is divided into three frames: The box on the left hand side shows the categories to which the various modules belong. When activated, this category selection is highlighted by a broad white frame. The selected, active category is color-highlighted. The corresponding modules of the active category are listed in a white-framed box on the right hand side of the window. At the bottom, find the buttons for 'Help' and 'Quit'.

After the first start of the YaST2 Control center, the uppermost category 'Software' is automatically selected. Change between categories using (↑) and (↓). Start a module belonging to the selected category by pressing (→). The module selection then appears highlighted by a broad white line. Select a module using (↑) or (↓). Scroll through the module selection by pressing either key continuously. When a module is selected, the title is color-highlighted. A short text describing this module is displayed in the bottom part of the window.

Start the desired module by pressing (Enter) when it is selected. Different buttons or selection fields of the module contain a differently-colored letter (yellow with the standard settings). The combination (Alt) + (yellow letter) selects the corresponding button directly.

Leave the YaST2 Control Center either using the 'Quit' button at the bottom part of the window or by choosing the 'Quit' menu item in the category selection and pressing (Enter).

Restriction of Key Combinations

It is possible that the (Alt) combinations in YaST2 do not work if system-wide (Alt) key combinations are set by a running X server. It is also possible that keys like (Alt) or (⇧ Shift) are captured for the terminal used.

Replacing (Alt) with (Esc): (Alt) shortcuts can be executed with (Esc) instead of (Alt). For example, (Esc) + (H) replaces (Alt) + (H).

Replacement of backward and forward navigation by (Ctrl) + (F) and (Ctrl) + (B): If the (Alt) and (⇧ Shift) combinations are occupied by the window manager or the terminal, the combinations (Ctrl) + (F) (forward) and (Ctrl) + (F) (backward) can be used instead.

Module Operation

In the following, it is assumed that the (Alt) key combinations are functional. Make appropriate substitutions or switch to a pure text console, if needed.

Navigation between buttons and selection lists (Tab) and (Alt) + (Tab) navigates back and forth between buttons and frames containing selection lists and among the frames.

Navigation in selection lists (↑) and (↓) always navigate among the single items within an activated frame containing a selection list. These can, for instance, be the single modules of a module group in the control center.

Checking radio buttons and check boxes The selection of buttons with empty square brackets (check boxes) or parentheses (radio buttons) can be done with the (Space) or (Enter) keys. The buttons at the bottom of the various modules or of the control center are activated with (Enter) when selected (colored green) or with the combination (Alt) + (yellow key) (cf. fig. D.2).
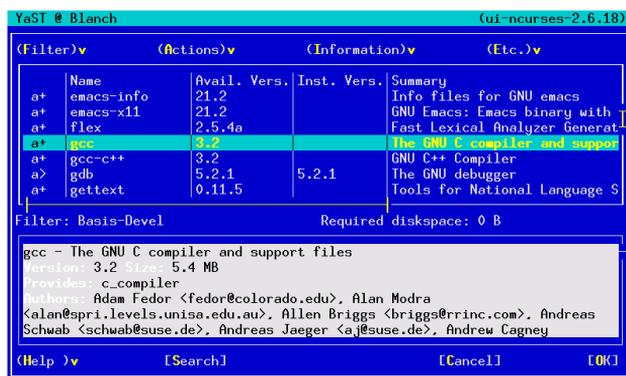
Figure D.2: *The Software Installation Module*

Invoking the Various Modules

Each YaST2 module can also be started directly. The modules can simply be started with yast ⟨*module name*⟩. The network module, for instance, is started with the command yast lan. Access a list of the names of the modules available on a system by running yast -l or yast --list.

LVM — The Logical Volume Manager

This text is based on the "LVM-HowTo" written by Heinz Mauelshagen.

Basics

LVM provides a very sophisticated way to handle disk space. You can build "Logical Volumes" by concatenating physical partitions. The partitions may be spread over different drives.

Example: You planned to use 600 MB for /home but actually you need 1 GB. LVM enables you to simply add another partition with 400 MB, for example, to the existing /home "on the fly". Without LVM, you would need to have a second partition with at least 1 GB and you would have to unmount the old /home, mount the new partition to /home, then copy all the data.

The handling of LVM is easiest when setting up a totally new logical volume. This can easily be done using YaST2. After building a file system on the new volume, use it like any other partition.

To add or remove space to a logical volume containing a running file system, use separate tools, like ext2resize, after growing the logical volume.

Further options: LVM provides the possibility of "striping". Another interesting feature is the "snapshot", which provides a non-persistent backup functionality.

Terms

Consider the terms used by LVM. Knowing these terms should help to understand the different YaST2 menus.

Physical Volume (PV)

A PV is a physical medium (e.g., `/dev/sda`), which is prepared for use by LVM. Therefore, some administrative data is added to it.

Physical Extent (PE)

PEs are like big blocks. A PV is subdivided into PEs. The default PE size is 4 MB.

Volume Group (VG)

A volume group contains a number of PEs, provided by one or more PVs. Do not worry about the PEs at this point, because you only have to tell LVM which PVs to use.

Logical Volume (LV)

Logical volumes are roughly the same as "partitions" and the Linux kernel does not differentiate between a regular partition and a logical volume. You can build any kind of supported file system on the LV. The number of LVs is currently limited to 256. This amount is shared among all VGs. The size of an LV is limited by the PE size configured at the time of volume group creation. When using the default value (4 MB), each LV is limited to 256 GB in size. If you need a bigger LV, choose another PE size. For example, LVs containing 1 Terabyte can be achieved by using a PE size of 16 MB.

Accessing the LVs

Assuming you have two PVs, `/dev/sda` and `/dev/sdb`, and a VG named "suse" containing the two PVs, the resulting LVs will be found at `/dev/suse/`. After creating a number of LVs named, for example, `test1`, `test2` `...testn`, find a hierarchy like

`/dev/suse/test1 /dev/suse/test2 .`
`.`
`.`
`/dev/suse/testn`

Mount these volumes like regular partitions.

Further Documentation

Assuming you have a running SuSE Linux Enterprise Server, find some documentation, including the LVM FAQ, at:

`/usr/share/doc/packages/lvm`

The LVM-HOWTO can be found at:

`/usr/share/doc/howto/en/LVM-HOWTO.gz`

The HTML Version can be accessed via:

`/usr/share/doc/howto/en/html/LVM-HOWTO.html`

Configuration with YaST2

Start the LVM configuration by selecting 'LVM...' in the 'Expert Partitioner' dialog.

The Partitioner

First, a dialog opens in which to change the partitioning of your hard disk. Here, remove or change current partitions as well as create new ones.
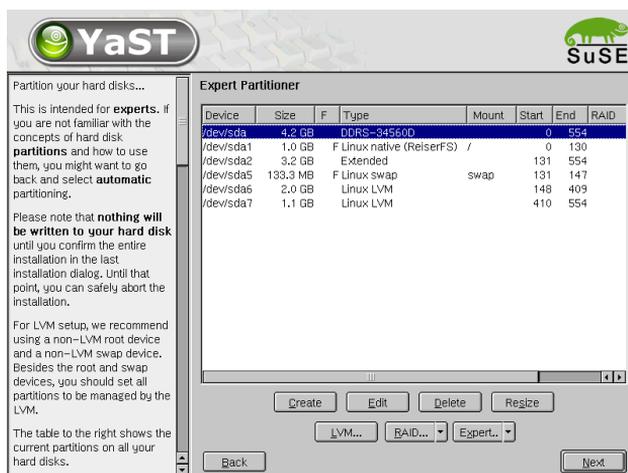
Figure E.1: YaST2: LVM Partitioner

As SuSE Linux Enterprise Server for S/390 and zSeries supports only two partition types ("Linux swap" and "Linux native"), a partition to use for

LVM should have the partition type "Linux native". The LVM tools recognize a PV automatically on boot regardless of the partition type.
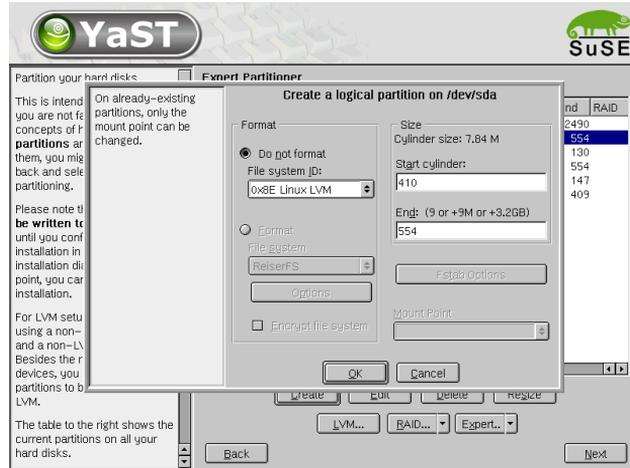
Figure E.2: YaST2: Creating an LVM partition

> **Note**
>
> If a valid LVM configuration already exists on your system, it will automatically be applied at the start of the LVM configuration. If this configuration is activated, no disks containing a partition belonging to an activated volume group can be repartitioned. The Linux kernel will refuse to detect the modified partitioning of a hard disk as long as even a single partition on this drive is used. Of course, repartitioning the disks not aassigned to an LVM volume group is not a problem. If you already have a valid LVM configuration on your system, it is usually not necessary to repartition it. In this dialog, configure all the mount points not located on the LVM logical volumes. In YaST2, at least the root file system must be located on a normal partition. Select this partition from the list and define it as the root file system using 'Edit'. Due to the great degree of flexibility in LVM, it is recommended to assign all other file systems to LVM logical volumes. After setting the root partition, exit this dialog.
>
> **Note**

Setting up Physical Volumes

This dialog manages the LVM volume groups (often abbreviated to "VG"). If there is no volume group yet on your system, you will be prompted by a pop-up window to create one. "System" is the name suggested for the volume group where your SuSE Linux Enterprise Server system files are located. The physical extent size (often abbreviated to PE size) defines the maximum size of a physical and logical volume in this volume group. This value is usually set to four megabytes. This allows for a maximum size of 256 gigabytes for a physical and logical volume. You should, therefore, only increase the physical extent size (for example, to 8, 16, or 32 megabytes) if you need logical volumes larger than 256 gigabytes.
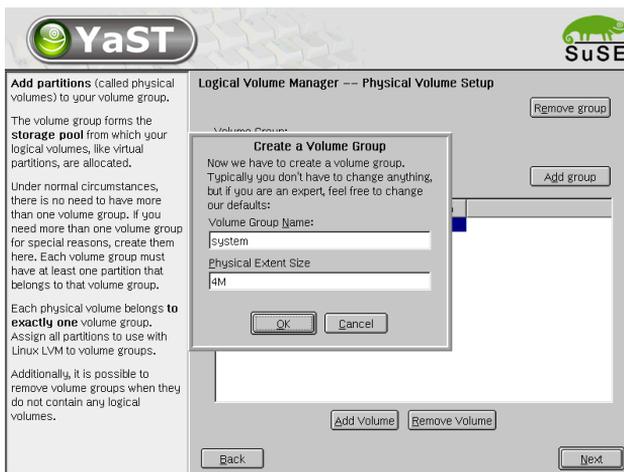
Figure E.3: YaST2: Creating a Volume Group

In the following dialog, all partitions are listed that either have "Linux LVM" or the "Linux native" types. If a partition is already assigned to a volume group, the name of the volume group is listed. Unassigned partitions bear the label "–".

The volume group currently edited can be modified in the selection box to the upper left. The buttons to the upper right enable creation of additional volume groups and deletion of existing volume groups. However, only volume groups without any additional partitions assigned to them can be removed. For a standard SuSE Linux Enterprise Server system, you do not need to create more than one volume group. A partition assigned to a volume group is also called a physical volume (often abbreviated to PV). To add
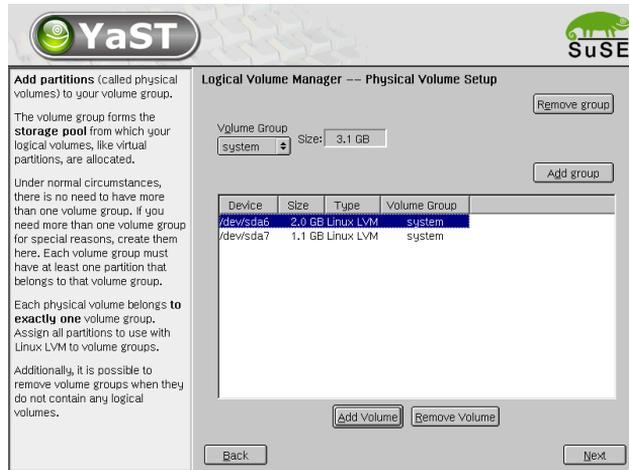
Figure E.4: YaST2: Overview of Partitions

a previously unassigned partition to the selected volume group, first select the partition then click 'Add volume' below the selection list. This allows the name of the volume group to be entered next to the partition selected. Assign all partitions intended for LVM to a volume group. Otherwise, the space on the partition will remain unused. Before exiting the dialog, assign at least one physical volume to each volume group.

Logical Volumes

This dialog manages the logical volumes (often just abbreviated to "LV").

Logical volumes are each assigned to a volume group and each have a certain size. Normally, a file system is generated on a logical volume (e.g., reiserfs, ext2), which is also designated a mount point. On an installed system, the files stored on this logical volume can then be found at this mount point. All the standard Linux partitions assigned a mount point, all swap partitions, and all already existing logical volumes are found in this list. If you have already configured LVM on your system, the available logical volumes should already be listed here. You will, however, still need to assign the appropriate mount point to these logical volumes. If you are configuring LVM on your system for the first time, there will not be any logical volumes yet in this screen and you will have to generate a logical volume for each mount point
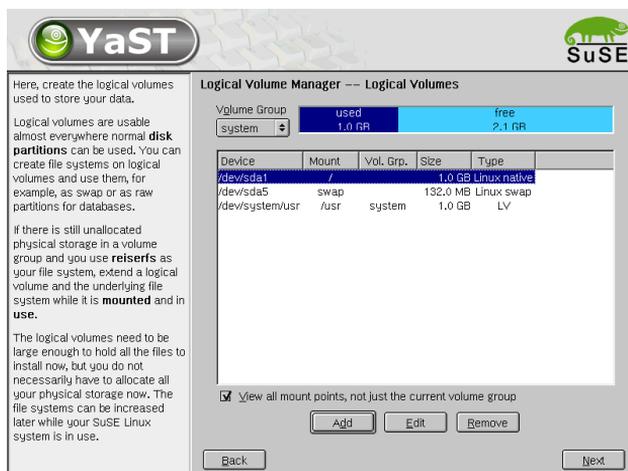
Figure E.5: YaST2: Management of the Logical Volumes

(with 'Add'), as well as define the size, the file system type (e. g., reiserfs or ext2), and the mount point (e. g., /var, /usr, /home).

If you have created several volume groups, switch between the different volume groups in the selection list to the upper left. The new logical volumes are all located in the volume group shown at the upper left. After creating all the logical volumes as required, the LVM configuration is complete. Then exit the dialog and continue to software selection if you are currently in the process of installation.

┌─ **Caution** ───

Implementing the LVM is also associated with increased risk factors such as data loss. Possible dangers are application crashes, power outages, and faulty commands. Secure your data before using LVM and before reconfiguring volumes. Do not work without backups.

─────────────────────────────────── **Caution** ─┘
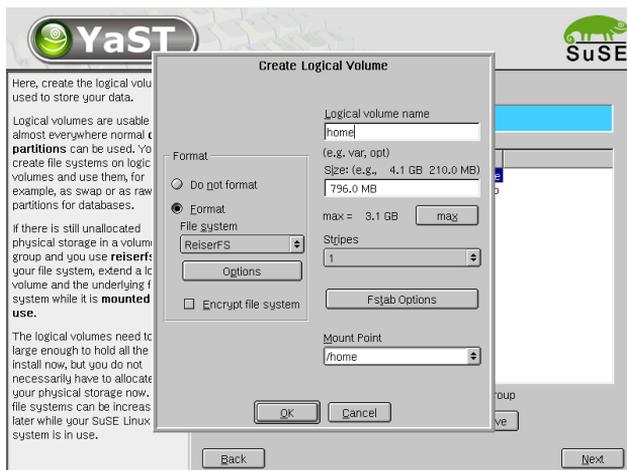
Figure E.6: *YaST2: Creating Logical Volumes*

File Systems in Linux

Linux supports a number of different file systems. This chapter presents a brief overview of the most popular Linux file systems, elaborating on their design concept, advantages, and fields of application. Some additional information on LFS ("Large File Support") in Linux is also provided.

Glossary

metadata A file system internal data structure that assures all the data on disk is properly organized and accessible. Essentially, it is "data about the data." Almost every file system has its own structure of metadata, which is partly why the file systems show different performance characteristics. It is of major importance to maintain metadata intact, because otherwise the whole data on the file system could become inaccessible.

inode Inodes contain various information about a file, including size, number of links, date, and time of creation, modification, and access, as well as pointers to the disk blocks where the file contents are actually stored.

journal In the context of a file system, a journal is an on-disk structure containing a kind of log where the file system stores what it is about to change in the file system's metadata. "Journaling" greatly reduces the recovery time of a Linux system since it obsoletes the lengthy search process that checks the whole file system at system startup. Instead, only the journal is replayed.

Major File Systems in Linux

Unlike two or three years ago, choosing a file system for a Linux system is no longer a matter of a few seconds ("Ext2 or ReiserFS?"). Kernels starting from 2.4 offer a variety of file systems from which to choose. The following is an overview of how those file systems basically work and which advantages they offer.

It is very important to bear in mind that there may be no file system that best suits all kinds of applications. Each file system has its particular strengths and weaknesses, which have to be taken into account. Even the most sophisticated file system cannot substitute for a reasonable backup strategy, however.

The terms "data integrity" or "data consistency", when used in this chapter, do not refer to the consistency of the user space data (the data your application writes to its files). Whether this data is consistent or not must be controlled by the application itself.

Ext2

The origins of Ext2 go back to the early days of Linux history. Its predecessor, the Extended File System, was implemented in April 1992 and integrated in Linux 0.96c. The Extended File System underwent a number of modifications and, as Ext2, became the most popular Linux file system for years. With the creation of journaling file systems and their astonishingly short recovery times, Ext2 became less important.

A brief summary of Ext2's strenghts might help you to understand why it was — and in some areas still is — the favorite Linux file system of many a Linux user.

Solidity Being quite an "old-timer", Ext2 underwent many improvements and was heavily tested. This may be the reason why people often refer to it as "rock–solid". After a system outage when the file system could not be cleanly unmounted, e2fsck starts to analyze the file system data. Metadata is brought into a consistent state and pending files or data blocks are written to a designated directory (called lost+found). In contrast to journaling file systems, e2fsck analyzes the whole file system and not just the recently modified bits of metadata. This takes significantly longer than checking the log data of a journaling file system. Depending on file system size, this procedure can take half an hour or more. Therefore, you would not choose Ext2 for any server that needs

to be highly available. Yet, as Ext2 does not maintain a journal and uses significantly less memory, it is sometimes faster than other file systems.

Easy upgradability   The code Ext2 is the strong foundation on which Ext3 could become a highly-acclaimed next-generation file system. Its reliability and solidity were elegantly combined with the advantages of a journaling file system.

Ext3

Ext3 was designed by Stephen Tweedie. In contrast to all other "next-generation" file systems, Ext3 does not follow a completely new design principle. It is based on Ext2. These two file systems are very closely related to each other. An Ext3 file system can be easily built on top of an Ext2 file system. The most important difference between Ext2 and Ext3 is that Ext3 supports journaling.

Summed up, Ext3 has three major advantages to offer:

Easy and highly reliable file system upgrades from Ext2   As Ext3 is based on the Ext2 code and shares its on-disk format as well as its metadata format, upgrades from Ext2 to Ext3 are incredibly easy. Unlike transitions to other journaling file systems, such as ReiserFS, JFS, or XFS, which can be quite tedious (making backups of the whole file system and recreating it from scratch), a transition to Ext3 is a matter of minutes. It is also very safe, as the recreation of a whole file system from scratch might not work flawlessly. Considering the number of existing Ext2 systems that await an upgrade to a journaling file system, you can easily figure out why Ext3 might be of some importance to many system administrators. Downgrading from Ext3 to Ext2 is as easy as the upgrade. Just perform a clean unmount of the Ext3 file system and remount it as an Ext2 file system.

Reliability and performance   Other journaling file systems follow the "metadata-only" journaling approach. This means your metadata will always be kept in a consistent state but the same cannot be automatically guaranteed for the file system data itself. Ext3 is designed to take care of both metadata and data. The degree of "care" can be customized. Enabling Ext3 in the `data=journal` mode offers maximum security (i.e., data integrity), but can slow down the system as both metadata and data are journaled. A relatively new approach is to use

the `data=ordered` mode, which ensures both data and metadata integrity, but uses journaling only for metadata. The file system driver collects all data blocks that correspond to one metadata update. These blocks are grouped as a "transaction" and will be written to disk before the metadata is updated. As a result, consistency is achieved for metadata and data without sacrificing performance. A third option to use is `data=writeback` which allows data to be written into the main file system after its metadata have been committed to the journal. This option is often considered the best in performance. It can, however, allow old data to reappear in files after crash and recovery while internal file system integrity is maintained. Unless you specify something else, Ext3 is run with the `data=ordered` default.

ReiserFS

Officially one of the key features of the 2.4 kernel release, ReiserFS has been available as a kernel patch for 2.2.x SuSE kernels since SuSE Linux Enterprise Server version 6.4. ReiserFS was designed by Hans Reiser and the Namesys development team. ReiserFS has proven to be a powerful alternative to the old Ext2. Its key assets are better disk space utilization, better disk access performance, and faster crash recovery. However, there is a minor drawback: ReiserFS pays great care to metadata but not to the data itself. Future generations of ReiserFS will include data journaling (both metadata and actual data are written to the journal) as well as ordered writes.

ReiserFS's strengths, in more detail, are:

Better disk space utilization   In ReiserFS, all data is organized in a structure called B*-balanced tree. The tree structure contributes to better disk space utilization as small files can be stored directly in the B*tree leaf nodes instead of being stored elsewhere and just maintaining a pointer to the actual disk location. In addition to that, storage is not allocated in chunks of 1 or 4 kB, but in portions of the exact size that is needed. Another benefit lies in the dynamic allocation of inodes. This keeps the file system more flexible than traditional file systems, like Ext2, where the inode density has to be specified at file system creation time.

Better disk access performance   For small files, you will often find that both file data and "stat_data" (inode) information are stored next to each other. They can be read with a single disk IO operation, meaning that only one access to disk is required to retrieve all the information needed.

Fast crash recovery   Using a journal to keep track of recent metadata changes makes a file system check a matter of seconds, even for huge file systems.

JFS

JFS, the "Journaling File System" was developed by IBM. The first beta version of the JFS Linux port reached the Linux community in the summer of 2000. Version 1.0.0 was released in 2001. JFS is tailored to suit the needs of high throughput server environments where performance is the ultimate goal. Being a full 64-bit file system, JFS supports both large files and partitions which is another reason for its use in server environments.

A closer look at JFS shows why this file system might prove a good choice for your Linux server:

Efficient journaling   JFS follows a "metadata only" approach like ReiserFS. Instead of an extensive check, only metadata changes generated by recent file system activity get checked, which saves a great amount of time in recovery. Concurrent operations requiring multiple concurrent log entries can be combined into one group commit, greatly reducing performance loss of the file system through multiple write operations.

Efficient directory organization   JFS holds to different directory organizations. For small directories, it allows the directory's content to be stored directly into its inode. For larger directories, it uses B$^+$trees, which greatly facilitate directory management.

Better space usage through dynamic inode allocation   For Ext2, you have to define the inode density in advance (the space occupied by management information), which restricted the maximum number of files or directories of your file system. JFS spares you these considerations — it dynamically allocates inode space and frees it when it is no longer needed.

Some Other Supported File Systems

Table F.1 on the following page summarizes some other file systems supported by Linux. They are supported mainly to ensure compatibility and interchange of data with different kinds of media or foreign operating systems.

cramfs	*Compressed ROM file system*: A compressed read-only file system for ROMs.
hpfs	*High Performance File System*: the IBM OS/2 standard file system — only supported in read-only mode.
iso9660	Standard file system on CD-ROMs.
minix	This file system originated from academic projects on operating systems and was the first file system used in Linux. Nowadays, it is used as a file system for floppy disks.
msdos	*fat*, the file system originally used by DOS, is today used by various operating systems.
ncpfs	file system for mounting Novell volumes over networks.
nfs	*Network File System*: Here, data can be stored on any machine in a network and access may be granted via a network.
smbfs	*Server Message Block*: used by products such as Windows to enable files to be accessed over a network.
sysv	Used on SCO UNIX, Xenix, and Coherent (commercial UNIX systems for PCs).
ufs	Used by BSD, SunOS, and NeXTstep. Only supported in *read-only* mode.
umsdos	*UNIX on MSDOS*: applied on top of a normal fat file system. Achieves UNIX functionality (permissions, links, long file names) by creating special files.
vfat	*Virtual FAT*: extension of the fat file system (supports long file names).
ntfs	*Windows NT file system*, read-only.

Table F.1: *File System Types in Linux*

Access Control Lists

Up to now one of the definite advantages of commercial UNIX systems was support for Access Control Lists (ACLs).

The traditional UNIX file object permission model allows read, write and execute permissions to be assigned to the owner, the owning group, and to the "rest of the world" (i.e., those users who are neither the owner nor members of the owning group). This scheme is sufficient in the majority of cases;

several workarounds to stretch its limits have been used by system administrators for years. Advanced applications nevertheless occasionally require a more flexible scheme.

In situations where the traditional file object permission model is no longer sufficient, Access Control Lists come to the rescue. They allow permissions to be defined for individual users and groups other than the owner or the owning group.

Access Control Lists are a feature of the Linux kernel, currently supported on ReiserFS, Ext2, Ext3, and XFS. With Access Control Lists more sophisticated scenarios can be realized directly, with no need to implement complex permission models in applications.

A popular example of where Access Control Lists marvel is replacing a Windows server by a Linux server. Often some of the workstations continue to run Windows. The Linux machine offers file and print services to the Windows clients via Samba. Because Samba supports Access Control Lists, permissions of users can be set on the Linux server, but also in the Windows GUI (Windows NT or later only). With the winbindd deamon it is even possible to define permissions for users who only exist in a Windows domain, but who have no accounts on the Linux server.

On the Linux server, Access Control Lists can be manipulated using the getfacl and setfacl utilities.

For further information on these commands refer to the man page for acl (man 5 acl), the man page for getfacl (man 1 getfacl) and the man page for setfacl (man 1 setfacl). Another useful source of information is the POSIX 1003.1e draft 17 specification, which can be obtained from http://wt.xpilot.org/publications/posix.1e/.

Extended Attributes

File objects have always been associated with a set of attributes like the file owner and owning group, permissions, and time stamps. The attributes that are available are predefined by the file system. Maybe you sometimes would have liked to attach some additional piece of information to a file, like the file's character encoding or a short description. With extended attributes this is finally possible.

Extended attributes (EAs) are name/value pairs associated permanently with files and directories, similar to the environment strings associated with a process. They were first introduced to Linux when support for the XFS file system was added. Meanwhile they are not only supported on XFS, but also on

ReiserFS, Ext2, Ext3 and JFS. These file systems also use extended attributes internally to implement Access Control Lists.

Extended attributes are not seen by any of the standard file operations. They are manipulated using special system calls. The getfattr and setfattr utulities can be used to retrieve and set extended attributes.

Further information on these commands can be found in the man page for attr (man 5 attr), the man page for getfattr (man 1 getfattr) and the man page for setfattr (man 1 setfattr). Table F.2 presents an overview of the file system limits for extended attributes.

File System	per EA	all EAs
ReiserFS	64kB	unlimited
Ext2/Ext3	File system block size (1/2/4/8kB)	File system block size (1/2/4/8kB)
XFS	64kB	unlimited
JFS	64kB	unlimited

Table F.2: File System Limitations for Extended Attributes

Large File Support in Linux

Originally, Linux supported a maximum file size of 2 GB. This was enough before the advent of multi media, and as long as no one tried to manipulate huge databases on Linux. Becoming more and more important for server computing, the kernel and C library were modified to support file sizes larger than 2 GB when using a new set of interfaces that applications must utilize. Nowadays, (almost) all major file systems offer LFS support allowing you to perform high-end computing.

Table F.3 on the facing page offers an overview of the current limitations of Linux files and file systems for Kernel 2.4. Hopefully the figures below will become outdated, with future kernel releases offering support for even larger files and file systems.

File System	File Size Limit	File System Size (Kernel Limit)
Ext2 or Ext3 (1 kB block size)	2^{34} Bytes (~16 GB)	2^{42} Bytes (4 TB)
Ext2 or Ext3 (2 kB block size)	2^{38} Bytes	2^{43} Bytes (8 TB)
Ext2 or Ext3 (4 kB block size)	2048 GB (2 TB)	16384 GB (16 TB)
Ext2 or Ext3 (8 kB block size) (Systems with 8 kB pages only (like Alpha))	65568 GB (~64 TB)	32768 GB (32 TB)
ReiserFS 3.5	4 GB	16384 GB (16 TB)
ReiserFS 3.6 (as in Linux 2.4)	2^{60} Bytes (1 EB)	16384 GB (16 TB)
XFS	2^{63} Bytes (8 EB)	2048 GB (2 TB) (Linux kernel limitation)
JFS (512 Bytes block size)	4194304 GB (4 PB)	512 TB
JFS (4 kB block size)	33554432 GB (32 PB)	4 PB

Table F.3: Maximum Sizes of File Systems

For More Information

Each of the file system projects described above maintains its own home page where you can find mailing list information as well as further documentation and FAQs.

```
http://e2fsprogs.sourceforge.net/
http://www.zipworld.com.au/~akpm/linux/ext3/
http://www.namesys.com/
http://oss.software.ibm.com/developerworks/opensource/jfs/
http://oss.sgi.com/projects/xfs/
```

A comprehensive multipart tutorial on Linux file systems can be found at *IBM developerWorks*:
http://www-106.ibm.com/developerworks/library/l-fs.html

For a comparison of the different journaling file systems in Linux, look at Juan I. Santos Florido's article at *Linuxgazette*: http://www.linuxgazette.com/issue55/florido.html.

Those interested in an in-depth analysis of LFS in Linux should try Andreas Jaeger's LFS site: http://www.suse.de/~aj/linux_lfs.html.

Support Services

SuSE Support Services offer a complete range of support services in connection with Linux. No matter whether you have questions on the installation of SuSE Linux Enterprise Server products or need to tackle an obstacle in your mission-critical IT solution, our service models adapt to your specific requirements.

For more information in the SuSE support services please contact one of the following SuSE offices or their website:

SuSE Germany
SuSE Linux AG
Business Support
Deutschherrnstr. 15-19
D-90429 Nürnberg
Phone: +49-911-74053-2330
Fax: +49-911-74053-489
bsupport@suse.de
http://www.suse.de/en/services/support

SuSE UK
SuSE Linux Ltd.
The Kinetic Centre, Theobald Street
Borehamwood, Herts. WD6 4PJ
Phone: +44-20-8387-4086
solutions@suse.co.uk
http://www.suse.co.uk

SuSE USA
SuSE Inc.
318 Harrison, #301
Oakland, CA 94607
Phone: (510) 628 3386
bsupport@suse.com
http://www.suse.com

Maintenance

Benefit from online access to the *SuSE Linux Maintenance Web* and ensure that
your system is always up-to-date and in a secure and stable state. Mainte-
nance for one year is already included in SuSE Linux Enterprise Server or
your SuSE Linux Business Solution and can be extended to further years.

Register your product at the following URL to get access to the Maintenance
Web:

http://support.suse.de/en/register/

After registration you will receive an e-mail describing further proceedings.
Once you have a login, you can directly access the *SuSE Linux Maintenance
Web* at:

http://support.suse.de/psdb/

Index

Article number: HB2111-3INST-INT

4 038564 007681