SuSE Linux Enterprise Server 8

for IBM iSeries and IBM pSeries

Installation

Edition 2004

Please direct suggestions and comments to documentation@suse.de

Authors: Dennis Geider, Thomas Fehr, Matthias Frühauf, Olaf Hering, Jana Jaeger, Peter Poeml, Martin Sommer, Marc Rührschneck
Translators: Rebecca Ellis, Olaf Niepolt
Editors: Antje Faber, Dennis Geider, Roland Haidl, Jana Jaeger, Edith Parzefall, Peter Reinhart, Michael Ranweiler, Marc Rührschneck, Thomas Schraitle, Martin Sommer, Rebecca Walter
Layout: Manuela Piotrowski, Thomas Schraitle
Setting: LaTeX

This book has been printed on 100 % chlorine-free bleached paper.

Contents

Introduction

About this document

This book describes the proceedings in installing SuSE Linux Enterprise Server on IBM iSeries and pSeries systems. You will be provided with all information needed to prepare an installation on the OS/400 side as well as the installation procedure of SuSE Linux Enterprise Server itself.

Wherever possible and appropriate we will provide links to more specific and most up-to-date documentation both on the net or as part of your installed Linux system.

Detailed information on Linux device drivers is excellently covered by IBM (online) documentation. Therefore, refer to the sources we will point to.

Structure of this document

Basically, this installation guide is divided in two parts.

Preparation This part will inform you about the requirements both on the hardware and the software side which need to be met in order to successfully install a SuSE Linux Enterprise Server on your machine. You will be informed which basic preparations needed to be performed on the 0S/400 side. The preparation on the firmware side on a pSeries system is also explained in detail.

Installing SuSE Linux Enterprise Server on IBM iSeries and pSeries 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 will also be provided.

Additionally, we will provide an appendix containing three basic chapters on handling YaST2 in textmode, LVM (Logical Volume Manager) and an overview of Linux filesystems.

Who Should Read this Document

We have made several assumptions concerning your background knowledge when designing this document.

- You are familiar with the aspects of OS/400 and pSeries firmware terminology.
- You have a good knowledge of the iSeries and pSeries devices attached to your system especially its network environment.
- You have a basic understanding of handling a Linux or Unix system.

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"
`newbie@earth:~ > ls`	enter the command `ls` in the shell of the user `newbie` in his home directory on the host "Earth"

Table 1: *continued overleaf...*

`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 Linux for iSeries and pSeries 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 iSeries and pSeries possible.

Nuremberg, 26th January 2004

Your SuSE team

Installation Overview

Before installing SuSE Linux Enterprise Server for *iSeries and pSeries*, make sure your system meets the basic system requirements. These requirements include the hardware and the software.

This part of the manual provides information about:

- Hardware requirements to install and run SuSE Linux Enterprise Server for *iSeries and pSeries*
- Software requirements to install and run SuSE Linux Enterprise Server for *iSeries and pSeries*
- Necessary steps to prepare for an installation on an iSeries system. This includes a description of a suitable telnet client, needed to communicate with the iSeries system during installation.
- Necessary steps to prepare for an installation on a pSeries system

This chapter will help with preparation and installation procedures to get your iSeries or pSeries system running a SuSE Linux Enterprise Server. Many chapters and sections in this manual are close with a “Where Next” showing you next steps to complete for the installation.

Installing on an IBM iSeries System

1. Check hardware requirements (Section *Hardware Requirements* on page 7)
2. Check software requirements (Section *Software Requirements* on page 8)

3. Prepare a VNC client for the installation (Section *Preparing a VNC Client for the Installation* on page 11)

4. Prepare the system on OS/400 side (Section *Preparing for Installation on IBM iSeries* on page 14)

5. Set up a client to access the iSeries console during the installation (Section *Preparing a Client to Control the Installation* on page 24)

6. IPL the kernel from OS/400 side (Section *IPL: Varying the NWSD On (WRKCFGSTS *NWS)* on page 32)

7. Install the software and do a basic network configuration (section *Installation with YaST2* on page 41)

8. Shut down the installed SuSE Linux Enterprise Server, configure the boot options on the OS/400 side, and restart Linux (Section *Booting Linux on an iSeries System* on page 59)

Installing on an IBM pSeries System

1. Check hardware requirements (Section *Hardware Requirements* on the facing page)

2. Check software requirements (Section *Software Requirements* on page 8)

3. Prepare the system to boot from the CD-ROM drive or network (Section *Preparing for Installation on IBM pSeries* on page 34)

4. Start the installation from CD-ROM or network (Section *Preparing for Installation on IBM pSeries* on page 34)

5. Install the software and do a basic network configuration (Section *Installation with YaST2* on page 41)

Requirements

Hardware Requirements

iSeries

At press time, we can support the following systems:

IBM iSeries 270
IBM iSeries 820
IBM iSeries 830
IBM iSeries 840
IBM iSeries 890

Please check `http://linuxppc64.org/boxes.shtml` for a detailed list including the feature codes. This information is also available at `http://www-1.ibm.com/servers/eserver/iseries/linux/pdfs/guide.pdf`.

Up-to-date information on the hardware requirements can be found at `http://www.ibm.com/servers/eserver/iseries/linux/reqs.html`.

A standard installation requires at least 84 Megabytes of memory. If you plan to install via VNC, you need at least 96 Megabytes of memory.

The installation of a "Default System" requires about 1.8 Gigabytes of hard disk space.

pSeries

Machines, which run a ppc32 kernel: RS/6000 43P Model 133
RS/6000 B50
RS/6000 150

RS/6000 170
RS/6000 260
RS/6000 270
RS/6000 pSeries 640

Machines, which run a ppc64 kernel:

44p 170
44p 260
44p 270
pSeries 610
pSeries 620
pSeries 630
pSeries 640
pSeries 660
pSeries 670
pSeries 690

Up-to-date information on the hardware requirements can be found at
`http://oss.software.ibm.com/developerworks/opensource/linux/projects/ppc/models.php`
and
`http://sdb.suse.de/en/sdb/html/olh_ppc_machines.html`.

A standard installation requires at least 84 Megabytes of memory. If you plan to install via VNC, you need at least 96 Megabytes of memory.

The installation of a "Default System" requires about 1.8 Gigabytes of hard disk space.

Software Requirements

IBM iSeries

Up-to-date information on the software requirements can be found at
`http://www.ibm.com/servers/eserver/iseries/linux/reqs.html`

Before starting the installation, you should check, if your system is up-to-date. IBM provides a list of PTFs (`http://www.ibm.com/servers/eserver/iseries/linux/ptfs.html`) necessary to run SuSE Linux Enterprise Server on iSeries and pSeries.

IBM pSeries

There are no special software requirements to run SuSE Linux Enterprise Server on pSeries systems. Linux will run natively on these machines.

Preparing

Preparing a VNC Client for the Installation

Introduction to VNC

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.

What is VNC used for?

SuSE Linux Enterprise Server uses VNC to enable the use of the graphical user interface of the installation software, YaST2. It is needed when systems cannot provide a graphical console or when 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, you need to prepare a remote computer, which is connected via network to the system you want to install on.

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.

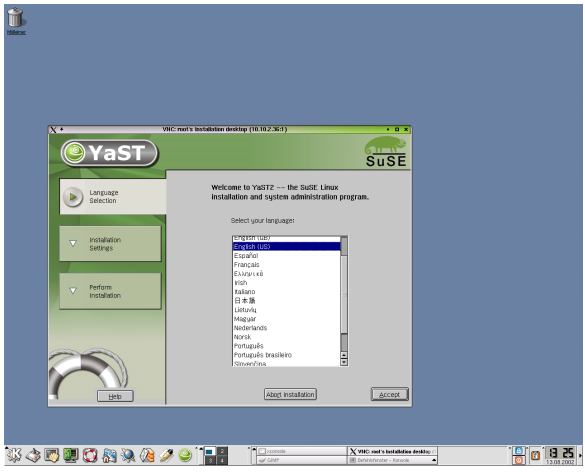

Figure 3.1: *YaST2 on a SuSE LinuxDesktop*

Preparing a Microsoft Windows Client

The VNC client for all Microsoft Windows 32bit 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 harddisk and unpack the file with the untgz32.exe, which can be found in the `dosutils/untgz` directory on CD 1.

```
C:\ tmp > untgz32.exe vnc-3.3.3r9_x86_win32.tgz
```

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

Initiating the Installation for VNC

The boot process for a VNC based installation differs slightly from a text based installation. The process as follows:

- Add the necessary kernel parameters to the boot prompt (see Table 3.1 on the next page)
- Boot the installation kernel
- If necessary, do the basic network configuration on the system console
- Start the vncclient on your client system and start installing the software

vnc=1	This enables the installation via VNC
vnc_password=xyz	This sets the password to access the VNC installation to xyz. Your VNC client will ask you for this password before you can access the installation procedure. Setting this option automatically enables the installation via VNC.
dhcp=1	With this option the installation procedure will automatically do a network configuration via DHCP. You need a working DHCP server on your network.
hostip=192.168.10.1	Sets the IP address to 192.168.10.1
netmask=255.255.255.0	Sets the netmask to 255.255.255.0
gateway=192.168.10.254	Sets the default gateway to 192.168.10.254
netdevice=eth0	Sets the network device to eth0
insmod=pcnet32	The pcnet32 module will aotomatically be loaded

Table 3.1: *Kernel parameters for VNC installation*

Additional Kernel Parameters

To enable the VNC installation, you need to add at least one parameter to the kernel parameter line. This needs to be done before the kernel is started.

The following table (*Additional Kernel Parameters* on this page) shows the kernel parameters useful in connection with a VNC installation:

Setting hostip, netmask, gateway, netdevice and the insmod parameter enables an automatic configured network with the given parameters.

In this manual, the installation via VNC is used on those platforms that do not support a graphical console.

Where Next

Proceed with the next section, Section *Preparing for Installation on IBM iSeries* on the next page, to start preparing the installation on IBM iSeries systems.

Proceed with section, Section *Preparing for Installation on IBM pSeries* on page 34, to start preparing the installation on IBM pSeries systems.

Preparing for Installation on IBM iSeries

Before starting the installation of SuSE Linux Enterprise Server on your iSeries system you need to prepare your system on the OS/400 side. This section covers installation via the built-in CD-ROM drive and via network.

Tip

The instructions on the following pages are suitable for a iSeries system running IBM OS/400 V5R1. It will also work on OS/400 V5R2. For additional options available on OS/400 V5R2, read the *Linux in a guest partition* available on the IBM website `http://publib.boulder.ibm.com/iseries/v5r2/ic2924/info/rzalm/rzalmlinuxkickoff.htm`

Tip

Refer to the article `http://sdb.suse.de/en/sdb/html/bugs-sles8-ppc.html` in the SuSE support database for up-to-date information on known problems, bug fixes, and workarounds for this release.

This chapter has received valuable comments and suggestions from Christopher Abbey, James Strebbing, Jay S. Bryant and Brent Baude.

Sources of Information

To get started:

- iSeries main site:
  `http://www.ibm.com/servers/eserver/iseries/linux/`
- The IBM "Linux on iSeries" Redbook:
  `http://www.redbooks.ibm.com/pubs/pdfs/redbooks/sg246232.pdf`
- The iSeries on Linux system guide:
  `http://www.ibm.com/servers/eserver/iseries/linux/pdfs/guide.pdf`
- Linux in a guest partition:
  `http://publib.boulder.ibm.com/iseries/v5r2/ic2924/info/rzalm/rzalmlinuxkickoff.htm`
- IBM Linux on LPAR information:
  `http://publib.boulder.ibm.com/pubs/html/as400/v5r1/ic2924/info/rzalm/rzalmlinuxkickoff.htm`

- Mailing lists:
  - `linuxppc-iseries` (to subscribe, send mail with "subscribe linuxppc-iseries" in the body to `Majordomo@lists.linuxppc.org`)
  - `linuxppc64-dev` (to subscribe, send mail with "subscribe linuxppc64-dev" in the body to `Majordomo@lists.linuxppc.org`)

For the maintenance of the installed system, these links are interesting:

- Support Database (SDB):
  `http://sdb.suse.de/sdb/en/html/index.html`
  The SDB is an effective means to help customers with difficulties. If SuSE notices that a specific issue does cause considerable problems or confusion, they will provide an SDB article. The SDB can be searched using key words like PPC or iSeries.
- Download page:
  `http://www.suse.com/us/support/download/updates/index.html`
- Security announcement page:
  `http://www.suse.com/us/support/security/index.html`
  SuSE also runs two security-related mailing lists to which any interested party may subscribe:
  - `suse-security@suse.com` — General, Linux, and SuSE security discussion. All SuSE security announcements are sent to this list. To subscribe, send an e-mail to suse-security-subscribe@suse.com>
  - `suse-security-announce@suse.com` — SuSE's announce-only mailing list. Only SuSE's security announcements are sent to this list. To subscribe, send an e-mail to suse-security-announce-subscribe@suse.com

Necessary Steps on the OS/400 Side

The following section will help with configuration of your iSeries Linux system for SuSE Linux Enterprise Server.

For detailed information on how create partitions for Linux, use the following information sources:

- For OS/400 V5R1:
  http://publib.boulder.ibm.com/pubs/html/as400/v5r1/ic2924/info/rzalm/rzalmlinuxkickoff.htm
- For OS/400 V5R2:
  http://publib.boulder.ibm.com/iseries/v5r2/ic2924/info/rzalm/rzalmlinuxkickoff.htm

This section only shows the important steps for the configuration. Refer to the Redbook "Linux on iSeries" (SG24–6232–00) for detailed instructions (http://www.redbooks.ibm.com/pubs/pdfs/redbooks/sg246232.pdf).

To configure your system on the OS/400 side, you need an OS/400 signon with *SERVICE authority and access to SST. A DST password is also required to create the console user account. To be able to access OS/400, you need a 5250 terminal or a 5250 emulation package (e.g. TN5250 on Linux or the PCS or Client Access on Windows).

Partitioning the System

First, you need to create a new system partition for the SuSE Linux Enterprise Server.

Use option 5 'Work with System Partitions', option 5 'Create a new Partition' in the STRSST, or option 3 -> 'Work with Partition Configuration' -> 2 'Change Partition Processing Resources' if you are working with an existing partition.

```
                          Create New Partition
                                                          System:   SUSE1
Complete blanks, press Enter.

  Partition identifier and name  ... . . . . .  3__  GINGER__

  Number of available system processors  ... . :  1
  Number of partition processors ... . . . . .  1__
  Minimum / maximum number of processors ... . .  0__ / 1__
  Use shared processor pool  ... . . . . . . .  2  1=Yes, 2=No

  Size of available system main storage (MB) ..:  256      / 44
  Size of partition main storage (MB)  ... . . .  256_____
  Minimum / maximum size of main storage (MB) ..  0_______ / 752_____




F3=Exit   F9=Exclude limits   F10=Work with shared processor pool
F11=Display partition processing configuration   F12=Cancel
```

```
                  Change Partition Processing Resources
                                                          System:   SUSE1
Type changes, press Enter.

  Partition identifier and name  .... . . . . :  3    GINGER

  Current / available number of processors ...:  1   / 0
  New number of processors .... . . . . . . .  1__
  Use shared processor pool  .... . . . . . .  2  1=Yes, 2=No

  Current / available size of main storage (MB)  :  256   / 44
  New size of main storage (MB)  .... . . . . .  256_____









F3=Exit   F9=Include limits   F10=Work with shared processor pool
F11=Display partition processing configuration   F12=Cancel
```

Make sure to set “Minimum / maximum number of processors” and “Minimum / maximum size of main storage” (F9) within the indicated ranges to avoid a system IPL.

```
                    Change Partition Processing Resources
                                                          System:   SUSE1
Type changes, press Enter.

  Partition identifier and name  . . . . . :  3    GINGER

  Current / available number of processors . . :  1   / 0
  New number of processors . . . . . .  1__
  Minimum / maximum number of processors . . . .  0__ / 1__
  Use shared processor pool  . . . . . .  2  1=Yes, 2=No

  Current / available size of main storage (MB) :  256      / 44
  New size of main storage (MB)  . . . . . .  256_____
  Minimum / maximum size of main storage (MB)  ..  0_______ / 752_____









F3=Exit   F9=Exclude limits   F10=Work with shared processor pool
F11=Display partition processing configuration   F12=Cancel
```

Then, apply your changes using Enter. This will start the creation process for the new partition.

Hints on Partitioning — Processors, Memory, NWSDs, and LPARs

- Use STRSST (Start System Service Tools), select option 5 ('Work with System Partitions', option 3 ('Work with Partition Configuration') and assign the hosting partition for the guest -> F13.
- Memory:
  There is a range in the LPAR (Logical Partition) config to setup. It is the minimum and maximum memory that can be assigned to the LPAR without a primary IPL (Initial Program Load). More than the configured amount matters: They reserve hardware page table (HPT) space for the MAX amount. This results in a small loss (Linux will not see the entire configured memory). So the max size should not be much larger than you plan to use.

  Here is the formula to calculate memory available to the LPAR:

$$\text{Configured Main Store} - \frac{\text{Maximum Main Store In LPAR Config}}{64}$$

  The fraction is rounded to power of 2.

Example:
Current maximum size is 248. Take 1/64 of that figure as allocated to the HPT, 3.875 MB rounded to 4 MB.

- Minimum memory for installation:
  To do a textbased installation, you need to assign at least 84 Megabytes of memory to the LPAR.

  If you plan to install via VNC, you need to assign at least 96 Megabytes of memory to the LPAR.

Creating Network Server Storage Space

Next, you need to create storage space for the SuSE Linux Enterprise Server. This is done with the `CRTNWSSTG` command.

```
                    Create NWS Storage Space (CRTNWSSTG)

Type choices, press Enter.

Network server storage space.. > GINGER0___     Name
Size.. . . . . . . . . . . . .   9000_______    *CALC, 1-64000 megabytes
From storage space.. . . . . .   *NONE_____     Name, *NONE
Format.. . . . . . . . . . . .   *open_____     *NTFS, *FAT,*FAT32,*OPEN
Auxiliary storage pool ID  . . .   1_____         1-99
Text 'description' . . . . . .   ginger_root_disk_____________________

                                                                  Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this dis
F24=More keys
Creating NWS storage space GINGER0: 32 of 9000 megabytes complete.
```

Creating the Network Server Description

The description puts the several configuration steps together into one object.

```
                    Create Network Server Desc (CRTNWSD)

Type choices, press Enter.

Network server description . . .   ginger__      Name
Resource name  . . . . . . . . .   *NONE_____    Name, *NONE
Network server type  . . . . . .   *guest____    *WINDOWSNT, *GUEST

                                                                  Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this dis
F24=More keys
Parameter NWSD required.                                            +
```

Change only the following settings:

- NWSD name
- Network server type = *GUEST
- Partition = <enter your partition here>
- Code page = 437
- IPL source = *STMF
- IPL stream file = '/QOPT/SU810.001/ISERIES64'
  This is different for a network based installation (see Section *Creating a Network Installation source* on page 30)
- IPL parameters = 'vnc=1 vnc_password=suseinst'
- Text description = 'SuSE Enterprise Server'
- Online at IPL = *NO

```
                    Create Network Server Desc (CRTNWSD)

Type choices, press Enter.

Network server description ... > GINGER__       Name
Resource name  ... . . . . . .   *NONE_____     Name, *NONE
Network server type  ... . . . > *GUEST____     *WINDOWSNT, *GUEST
Online at IPL  ... . . . . . .   *NO           *YES, *NO
Vary on wait ... . . . . . . .   *NOWAIT       *NOWAIT, 1-15 minutes
Partition  ... . . . . . . . . > GINGER____     Name
Code page  ... . . . . . . . .   437____       *LNGVER,437,850,852,857
Server message queue ... . . .   *JOBLOG___    Name, *JOBLOG, *NONE
  Library  ... . . . . . . . .     __________  Name, *LIBL, *CURLIB
TCP/IP port configuration:       _
  Port ... . . . . . . . . . .   *NONE_____    *NONE, *INTERNAL, 1, 2,3
  Internet address ... . . . .   ______________________
  Subnet mask  ... . . . . . .   ______________________
  Maximum transmission unit ..   ___________   Number
            + for more values _
TCP/IP local host name ... . .   *NWSD_____________________________________
TCP/IP local domain name ... .   *SYS_____________________________________
TCP/IP name server system  ...   *SYS_____________
            + for more values    ________________
Restricted device resources ..   *NONE_____    Name, *NONE, *ALL...
            + for more values    __________
Synchronize date and time  ...   *TYPE         *TYPE, *YES, *NO
IPL source ... . . . . . . . . > *STMF__       *NWSSTG, *PANEL, *STMF,A
IPL stream file  ... . . . . . > '/QOPT/SU8100.001/ISERIES64'__________
IPL parameters ... . . . . . .   'vnc=1 vnc_password=suseinst'___
Authority  ... . . . . . . . .   *LIBCRTAUT    Name, *LIBCRTAUT,*CHANGE
Text 'description' ... . . . . > 'SuSE_Enterprise_Server'_______________

                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this dis
F24=More keys
```

You can add additional IPL parameters for VNC. Please refer to table *Additional Kernel Parameters* on page 13.

You can look up the IPL stream file via the `DSPLNK('QOPT')` command in the OS/400 command entry.

Linking the Network Storage Space to the Network Server Description (WRKNWSSTG)

You also need to assign the new created storage space to the description. Select the storage space from the list first.

```
                    Work with Network Server Storage Spaces
                                                              System:   SUSE1
Type options, press Enter.
  1=Create   3=Copy   4=Delete   5=Display   6=Print   10=Add link
  11=Remove link

                 Percent
Opt  Name        Used     Size   Server  Drive  Format    Access    ASP
__   __________
__   CURRY0         0     2000  CURRY       1   *OPEN     *UPDATE     1
__   CURRY1         0     2000  CURRY       2   *OPEN     *UPDATE     1
10   GINGER0        0     9000                  *OPEN                 1
__   TEST           0     1500                  *OPEN                 1









                                                                    Bottom
Parameters or command
===> ____________________________________________________________________________
F3=Exit   F4=Prompt   F5=Refresh   F6=Print list   F9=Retrieve
F11=Display text      F12=Cancel   F17=Position to
```

And then link it to the Network Server Description.

```
                     Add Network Server Storage Link (ADDNWSSTGL)

Type choices, press Enter.

Network server storage space . . > GINGER0       Name
Network server description . . . > GINGER        Name
Drive letter . . . . . . . . . .   *FIRSTAVAIL   K-Z
Dynamic storage link . . . . . .   *yes          *NO, *YES
Network server type  . . . . . .   *NWSD         Character value
Drive sequence number  . . . . .   *CALC         3-18, *CALC

                           Additional Parameters

Access . . . . . . . . . . . . .   *UPDATE       *UPDATE, *READ









                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this dis
F24=More keys
```

Configuring the Virtual Ethernet (System Tools)

To configure the virtual ethernet, use the STRSST (System Service Tools) again. Select option 5 'Work with system partitions', option 3 'Work with partition configuration' and then use F10 to change the virtual lan configuration.

```
                    Work with Virtual LAN Configuration
                                                          System:   SUSE1
Type options, press Enter.
  2=Change

    Par          --------------Virtual LAN Identifiers---------------
Opt ID  Name     0  1  2  3  4  5  6  7  8  9  10  11  12  13  14  15
 _    0 PRIMARY  1  .  .  .  .  .  .  .  .  .   .   .   .   .   .   .
 _    1 PEPPER   1  .  .  .  .  .  .  .  .  .   .   .   .   .   .   .
 _    2 CURRY    1  .  .  .  .  .  .  .  .  .   .   .   .   .   .   .
 _    3 GINGER   1  .  .  .  .  .  .  .  .  .   .   .   .   .   .   .









'1' Indicates LAN in use.  '.' Indicates LAN not in use.
F3=Exit                            F9=Show only partitions using Virtu
F11=Display communication options   F12=Cancel
```

To display all partitions (including the ones not yet connected), use F9.

```
                    Change Virtual LAN Configuration
                                                              System:   SUSE1
  Partition identifier . . . . . . . . . . . . . . . :  3
  Partition name . . . . . . . . . . . . . . . . . . :  GINGER

Type changes, press Enter.
  1=Yes  2=No

  -------------------Virtual LAN Identifiers-------------------
  0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
  1   2   2   2   2   2   2   2   2   2   2   2   2   2   2   2









F3=Exit   F12=Cancel
```

Where Next

Proceed with the next section, Section *Preparing a Client to Control the Installation* on the current page, to configure a virtual console to be able to control the system during the installation.

Preparing a Client to Control the Installation

This section is about configuring telnet access to the virtual console.

Supported Terminals

The terminal emulators that are supported are:

Linux	Linux console	fully supported
Linux	standard XTerm	fully supported
Linux	GNOME terminal	fully supported
Linux	KDE console	fully supported
Linux	screen	fully supported
Windows	PuTTY (free telnet and ssh client)	fully supported
Windows	telnet client shipping with Windows 98	not supported
Windows	telnet client shipping with Windows 2000	not recommended
Windows	telnet client shipping with Windows XP	not recommended
AIX	aix xterm	not supported

Note

Using a fully supported terminal emulator enables you to use all features of the installer including shortcut keys and colors.

Note

For an installation done from a Windows computer we strongly recommend you use the free telnet client PuTTY. Please see below for directions on its use and configuration.

The Windows 98 telnet client is not supported. The Windows 2000 and Windows XP telnet clients are usable, but can only display the installer in a limited way.

aixterm does not work very well together with Linux in general. A workaround on AIX is to start an xterm and set TERM=vt100 manually. Since this is not possible during the installation, though, it is advisable to use another platform for virtual console access to do the installation.

In the Linux terminal multiplexer screen it is recommended to change the background-color-erase setting, because otherwise the background will be black instead of the background color. This can be achieved by adding the line `defbceon` to `~/.screenrc` which changes the default. To change the setting for a single window only, the keystroke sequence Ctrl + A + ":bce" can be used.

PuTTY Setup

PuTTY is a free telnet and ssh client for Windows with very good capabilities. putty.exe can be found on CD 1 in the directory `dosutils`. It can also be obtained from `http://www.chiark.greenend.org.uk/~sgtatham/putty/latest/putty.exe`.

CD 1 is readable on any Windows computer. The program consists of one executable, and no installation is needed besides copying the file `dosutils\putty\putty.exe` to your hard disk.

The full documentation is to be found on its homepage: `http://www.chiark.greenend.org.uk/~sgtatham/putty/`

For ease of access, it is useful to create and save a session profile for the virtual console in the PuTTY configuration dialog. This is described below.

It is recommended to configure the following settings:

- A larger font (the font that PuTTY uses as default is tiny).
- Set the font to "Courier New" instead of "Courier". This allows for correct line drawings. The font "Fixedsys" can also be used.

- Set a larger window size, e.g. 40x100 (default is 24x80).
- Create a session profile named after the host with these settings for access to port 2301 of you iSeries.

Double-click the PuTTY icon. A dialog box appears (see Figure 3.2).

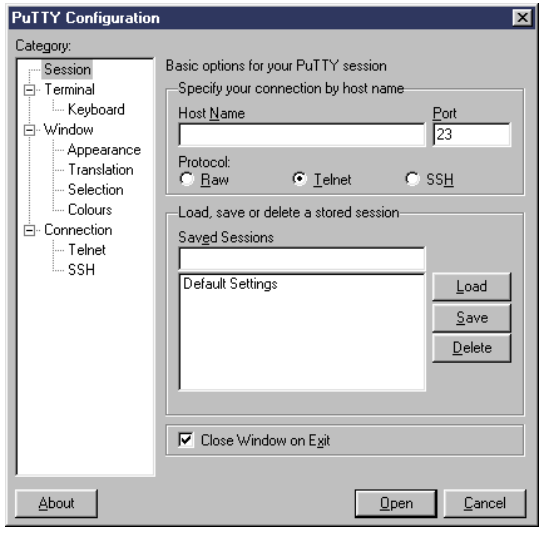

Figure 3.2: *First start of PuTTY*

1. Select 'Window' and change the window size to e.g. 40x100 (entering the values respectively into the 'Rows' and 'Columns' fields, see Figure 3.3).
2. Select 'Appearance' and click the 'Change' button (see Figure 3.4).
3. Choose a font (see Figure 3.5). "Courier new", size 10, is suggested. If this font is not available, "Fixedsys", size 12, can also be used.
4. Click the 'OK' button to save the font change. Note: It is possible that now a different font is shown, which seems to be a bug in PuTTY (see Figure 3.6).
5. Select 'Session' (upper left corner). The dialog box now appears as it looked initially.
6. To use PuTTY to login to the host named `iSeries`, fill in the 'Host name' field in the PuTTY dialog box (see Figure 3.7).

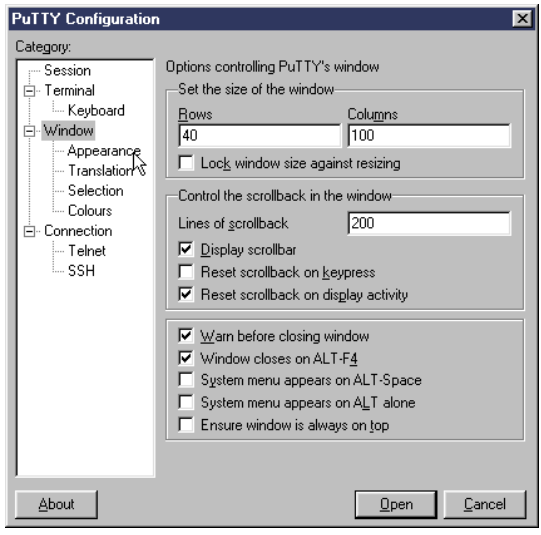

Figure 3.3: *Changing the size of the PuTTY window*

7. Enter `2301` into the 'Port' field.
8. Enter a descriptive name into the 'Saved Sessions' field, e.g. "iSeries virtual console".
9. Click the 'Save' button on the right side (see Figure 3.8).

After the session is configured, click on its name, and then on 'Open' to connect to the virtual console.

The Virtual Console Terminal Detection

It might help to know about an issue with telnet access to the iSeries virtual console, which is relevant for the installation as well as for virtual console access to an already installed system:

The OS/400 telnet server has the peculiarity that it allows multiple connects accessing the very same Linux virtual console.

To achieve optimized results for a given connected telnet client, Linux probes it. Specifically, it tries to determine

- whether more than one terminal is connected (this results in a warning),

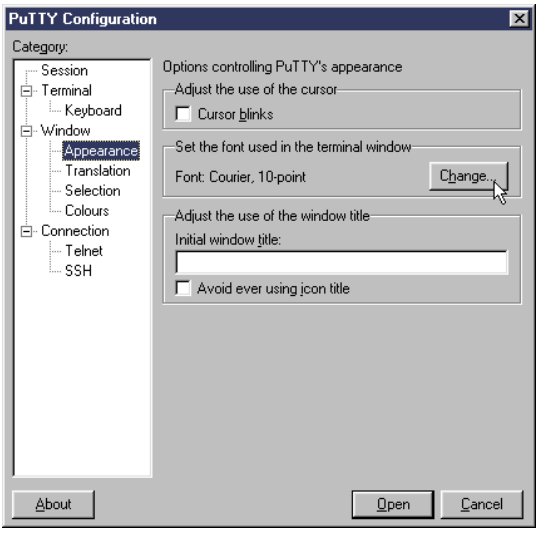

Figure 3.4: *Open the font dialog in PuTTY*

- the type of the terminal connected,
- whether it is a Windows-based terminal, and
- the screen size of the terminal.

This detection is done when the installer is started, as well as when a user logs in to the installed system over the virtual console.

Note

If more than one session is connected, all terminals will reply to the probing and the replies cannot be interpreted unambiguously. To achieve predictable results, multiple connections should not be used.

Note

Forcing Terminal Initialization

The terminal detection and initialization can also be forced manually (in the installed system) by entering the `initviocons` command.

Whenever the terminal is resized while being logged in over the virtual console, the `initviocons` command should be used to reinitialize the terminal in order

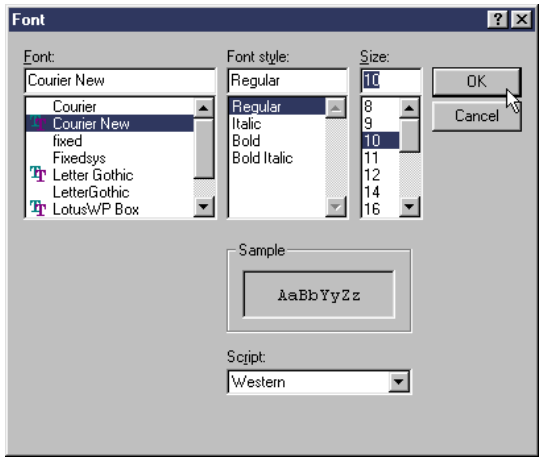

Figure 3.5: *Changing the PuTTY font*

to make Linux aware of the changed screen size, so all screen-based programs can work properly. However, since the TERM variable cannot be set retroactively in the environment of the user, it is recommended to simply sign on again after changing the telnet client.

Troubleshooting Terminal Problems

- In general, Ctrl + L is the command to redraw the screen, should this be needed.
- Linuxrc (the blue screen you get before YaST2 opens) shows wrong line drawings (it only uses hard-wired control codes). This is a known problem which does not indicate a problem on the telnet side.
- If you are using a terminal wherein a previous install has been cancelled, the terminal might be in an undefined state. This can happen e.g. after an immediate poweroff of the LPAR. If you experience problems it should help to reconnect in a fresh window.
- If the Backspace key does not work in the KDE "konsole", try either one of the 'Keyboard' settings (in the 'Settings') `xterm (XFree 3.x.x)` or `linux console`. It is possible that a different setting is needed to make the backspace key work after the system has been booted.

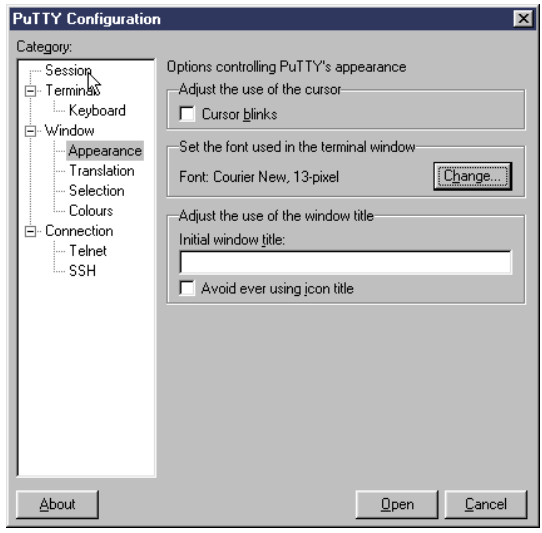

Figure 3.6: *The new PuTTY font*

- In the VT100 fallback mode, the Backspace key might not work. This obviously also depends on the configuration of the telnet client. Often you can work around by pressing (Ctrl) + (H) to achieve the backspace effect.
- If the (Enter) key does not respond as expected, use (Ctrl) + (J) instead.
- Generally, when the Backspace key does not work as expected, you could also try either one of these keys:
  - (Ctrl) + (H)
  - (Ctrl) + (?)
  - (Ctrl) + (Backspace)
  - (Delete) (used together with (←))
  - (Ctrl) + (D) (used together with (←))

Creating a Network Installation source

If you are planning on installing SuSE Linux Enterprise Server on many partitions it may be worth your time to create a network installation source. Having a

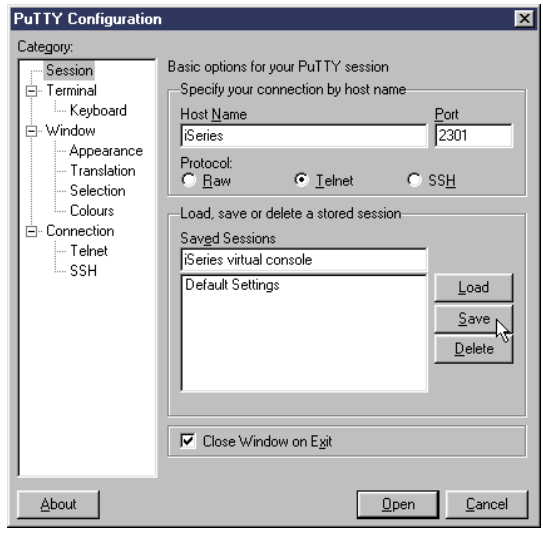

Figure 3.7: *Creating the Putty session*

network installation source has the advantage of allowing the user to install without having to change CDs and allows you to install onto multiple partitions simultaneously with ease.

To create the network share, copy the installation CDs into individual directories on a system that can act as a NFS server. For example on an existing SuSE Linux box you could copy each CD using the command

```
earth:/ #  cp -a /media/cdrom /suse-share/
```

then rename the cdrom directory that is created to “CD1”. I.E.

```
earth:/ #  mv /suse-share/cdrom /suse-share/CD1
```

Repeat the process for the remaining CDs. Then export the `/suse-share` directory via NFS.

For your installation source you will need to FTP the `ISERIES64` file to the OS/400 IFS using the following steps:

```
earth:/ #  ftp iseries
(login)
ftp> cd /kernels
ftp> bin
```

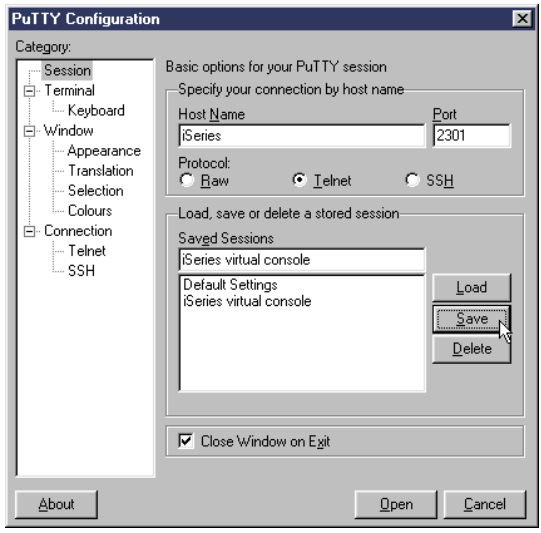

Figure 3.8: *Saving the Putty session*

```
ftp> put ISERIES64
ftp> bye
```

Now you can install using the steps outlined above with one change, your IPL Source will be as follows:

```
   IPL source . . . . . . . . . . . *STMF__ *SAME, *NWSSTG, *PANEL...
   IPL stream file . . . . . . . . . '/kernels/ISERIES64'_________________
   IPL parameters . . . . . . . . .'vnc=1 vnc_password=suseinst'________
```

You can add additional IPL parameters for VNC. Please refer to table *Additional Kernel Parameters* on page 13.

Where Next

Proceed with the next Section *IPL: Varying the NWSD On (WRKCFGSTS *NWS)* on the current page to finally IPL the installation kernel.

IPL: Varying the NWSD On (WRKCFGSTS *NWS)

Now, connect to the virtual console using the PuTTY session you just created, or by connecting the system via telnet on port 2301 using a supported terminal.

Finally, you can IPL the installation kernel by selection option 1 in front of the NWSD for the partition you wish to IPL. Please also watch the configured virtual console to see the kernel messages.

```
                        Work with Configuration Status          SUSE1
                                                        12/03/01  17:45:21
Position to  . . . . .   __________   Starting characters

Type options, press Enter.
  1=Vary on   2=Vary off   5=Work with job   8=Work with description
  9=Display mode status    13=Work with APPN status...

Opt  Description       Status                ----------Job-------------
__   CINST            VARIED OFF
__   CURRY            ACTIVE
1_   GINGER           VARIED OFF
__   PEPPER           VARIED OFF
__   PEPPERI          ACTIVE
__   TEST             VARIED OFF




                                                                   Bottom
Parameters or command
===> ____________________________________________________________________
F3=Exit   F4=Prompt   F12=Cancel   F23=More options   F24=More keys
```

Watch the virtal console. If you install via network or VNC, you will enter LinuxRC now:

1. Select your language for LinuxRC.
2. Select 'Kernel modules (hardware drivers)'.
3. Select 'Load ppc_iseries modules' and choose the network driver for your network card.
4. Select 'Back' and then 'Start installation / system'.
5. Select 'Start installation / system'.
6. Select the installation source. Three options are possible:

   (a) CD-ROM
   Select this option, if you are using the installation media shipped with your SuSE Linux Enterprise Server. Before the system starts accessing the CD-ROM drive, you will be asked for the basic network parameters. This is necessary to enable VNC.

(b) Network
Allows you to access the installation data e.g. from an NFS-shared directory. You will be asked for the basic network parameters for your system and then for the Hostname/IP address of the NFS server. Finally, enter the path to the NFS-shared installation data (e.g.`suse-share`).

(c) Harddisk
This is useful, if you have the installation data available on a separate harddisk. Enter the devicename with the partition and the path to the installation data in the following dialog (e.g. `/dev/sdb1` and `/suse`). Before the system starts accessing the harddisk drive, you will be asked for the basic network parameters. This is necessary to enable VNC.

Finally, a message will appear prompting you to start your VNC client on the remote system (See *IPL: Varying the NWSD On (WRKCFGSTS *NWS)* on this page).

```
starting VNC server...
a log can be found in /tmp/vncserver.log ...

***
***           You can connect to 192.168.0.154, display :1 now
***

(When YaST2 is finished, close your VNC viewer and return to this window.)
```

Output 1: *The System is ready to be maintained via VNC*

Start the VNC client with the parameters shown in the output (192.168.0.154:1 in our example) and enter the VNC password (“suseinst” in our example). The VNC graphical screen will appear and a few seconds later, YaST2.

Where Next

Proceed with Chapter *Installation with YaST2* on page 41 to start the actual installation of the software.

Preparing for Installation on IBM pSeries

This section describes necessary steps prior starting an installation of SuSE Linux Enterprise Server Enterprise Server on an pSeries system. It will cover the installation via the built-in CD-ROM drive and via a network.

Creating a Network Installation source

If you are planning on installing SuSE Linux Enterprise Server Enterprise Server on many partitions it may be worth your time to create a network installation source. Having a network installation source has the advantage of allowing the user to install without having to change CDs and allows you to install onto multiple systems simultaneously with ease.

To create the network share copy the installation CDs into individual directories on a system that can act as a NFS server. For example on an existing SuSE Linux box you could copy each CD using the command

```
earth:/ #  cp -a /media/cdrom /suse-share/
```

then rename the cdrom directory that is created to "CD1". I.E.

```
earth:/ #  mv /suse-share/cdrom /suse-share/CD1
```

Repeat the process for CD2. Then export the `/suse-share` directory via NFS.

IBM pSeries p670 and p690 Issues

The IBM p670 and p690 systems provide a partitioning functionality similar to the one available on IBM iSeries systems. This allows you to run up to 16 operating environments simultaneusly on one system. These operating environments are stored in *LPAR*s (Logical Partitions). One or more of those partitions could hold a SuSE Linux Enterprise Server environment.

To prepare an LPAR for SuSE Linux Enterprise Server you need to configure the system using the *HMC* shipped with the system. Check the *IBM @server pSeries 690 System Handbook* Redbook (SG24-7040-00) for details on that.

Note the following hints on the configuration:

- Giving more than 8 processors to the SuSE Linux Enterprise Server LPAR is not recommended, as the current Linux kernel doesn't scale very well on systems with more than 8 CPUs.
- For the installation mode, select 'SMS' as the boot mode for the partition.
- The *HMC* textmode Terminal used to control the installation is a VT320 clone. This could cause some strange behaviour on some applications. Whenever possible, use an xTerm to connect to the partition.

Hard Disk Space

First make sure you have enough disk space left to install SuSE Linux Enterprise Server Enterprise Server. Using a new hard disk for the installation is recommended. Refer to section *Hardware Requirements* on page 7 for the required hard disk space.

Making the System Bootable from CD-ROM Drive

Insert CD 1 into the CD-ROM drive and reboot the system. Next, enter the firmware of your pSeries system to enable booting from a CD-ROM drive. Normally, this is done by pressing F1 during the firmware checks at boot.

```
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000



     memory      keyboard     network      scsi      speaker
```

Press F1 when the SCSI devices are checked.

Next, you should see the 'Multiboot' dialog:

```
RS/6000 Firmware
Version SPX98285
(c) Copyright IBM Corp. 1997  All rights reserved.
-------------------------------------------------------------------------------
Multiboot

1  Select Software
2  Software Default
3  Select Install Device
4  Select Boot Devices
5  OK Prompt
6  Multiboot Startup <ON>









                                                                    .-------.
                                                                    |X=Exit|
                                                                    '------'


===>3
```

Select 3 to change the “Install Device”.

A list of valid devices appears.

```
Install Operating System

Device Device
Number Name
1        Diskette
2        SCSI Tape id=0 ( slot=50322f5a )
3        SCSI CD-ROM id=1 ( slot=50322f5a )
4        Ethernet  ( Integrated )
5        SysKonnect PCI FDDI Adapter  ( slot=4 )
6        Ethernet  ( slot=2 )
7                None









                                                          .------.
                                                          |X=Exit|
                                                          '------'

===>3
```

Select the appropriate CD-ROM drive (e.g., 3 in our example).

The system starts reading from the CD-ROM drive and displays the “Identstring”.

```
->1   SuSE Linux SLES-8 (PPC)<-
```

Booting from the CD-ROM Drive

After pressing ①, the yaboot screen will appear.

```
Config file read, 450 bytes


Welcome to SuSE Linux!

  Use  "install"     to boot the ppc64 kernel
  Use  "power3"      to boot the 32bit POWER3 kernel
  Use  "install32"   to boot the 32bit kernel

  You can pass the option "noinitrd"  to skip the installer.
  Example: install noinitrd root=/dev/sda4



Welcome to yaboot version 1.3.6.SuSE
boot:
* install                   power3                     install32

You can also type in custom image locations, in the form
{prom_path;}partno/path_to_image or {prom_path;}{partno}[start-end]
Example: hd:3,/vmlinux

boot:
```

Select 'install' from the menu or just hit Enter. To install on a POWER3 system, enter "power3" and press Enter. To install on a 32bit system, enter "install32" and press Enter.

If you plan to do a installation from a network source (see *Creating a Network Installation source* on page 35), you need to add "manual" in addition to the kernel (install, power3 or install32).

In case you plan to use VNC for the installation, add at least the "vnc=1" and "vnc_password=" parameter to the kernel (install, power3 or install32). Read more on VNC in Section *Preparing a VNC Client for the Installation* on page 11.

If the start of yaboot fails, you have to inititate the load manually:

- During the hardware test, hit F8 or 8. You will enter the firmware prompt.
- Check the available device aliases by listing them:

  ```
  0> devalias
  ```

- Enter the boot string including the devalias (`cdrom` in this example) To

boot a ppc64 kernel:

```
0> boot cdrom:,install
```

To boot a ppc32 kernel:

```
0> boot cdrom:,install32
```

To boot a power3 kernel:

```
0> boot cdrom:,install_power3_32
```

If you plan to do a installation from a network source (see *Creating a Network Installation source* on page 35), you need to add “manual” in addition to the kernel (install, power3 or install32).

In case you plan to use VNC for the installation, add at least the “vnc=1” and “vnc_password=” parameter to the kernel (install, power3 or install32). Read more on VNC in Section *Preparing a VNC Client for the Installation* on page 11.

The kernel will be loaded and started. Watch the kernel messages on the console.

If you decided to do a installation from a network source, LinuxRC will appear and the following steps are necessary:

1. Select your language for LinuxRC.
2. Select ‘Kernel modules (hardware drivers)’
3. Select ‘Load ppc_pseries modules’ and choose the network driver for your network card (e.g. “pcnet32”).
4. Select ‘Back’ and then ‘Start installation / system’.
5. Select ‘Start installation / system’.
6. Select ‘Network’, ‘NFS’ as your source medium.
7. Select your network device if prompted.
8. Enter the IP address and networking information when prompted.
9. Enter the IP of the system you created the NFS share on in Section *Creating a Network Installation source* on page 35.
10. Enter the path to the NFS share. I.E. "/suse-share/CD1" using the example from *Creating a Network Installation source* on page 35.

Finally, you will be asked for the terminal type you use.

```
What type of terminal do you have ?

  1) VT100
  2) VT102
  3) VT220
  4) X Terminal Emulator (xterm)
  5) X Terminal Emulator (xterm-vt220)
  6) X Terminal Emulator (xterm-sco)
  7) X Terminal Emulator (xterm-sun)
  8) Linux VGA or Framebuffer Console
  9) Other

Type the number of your choice and press Return: 1
```

YaST2 will be loaded and then started.

Where Next

Proceed with Section *Installation with YaST2* on the next page to start the installation of the software.

Installation with YaST2

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. If your cursor does not move, your mouse has not been automatically recognized by Linux. You will then need to use your keyboard. If YaST2 was started in textmode, you cannnot use the mouse. Refer to Chapter *YaST2 in Text Mode (ncurses)* on page 65 for details on the usage.

Selecting a Language

SuSE Linux Enterprise Server and YaST2 are adapted to use the language selected.

If your mouse does not work, navigate with the arrow keys to the desired language then press Tab repeatedly until the 'Next' is selected. Then press ↵.

Installation Mode

Here you can decide if you want to do a new installation, or just boot an already installed system (Figure 4.2 on page 43). This might be necessary in cases like a misconfigured bootloader.

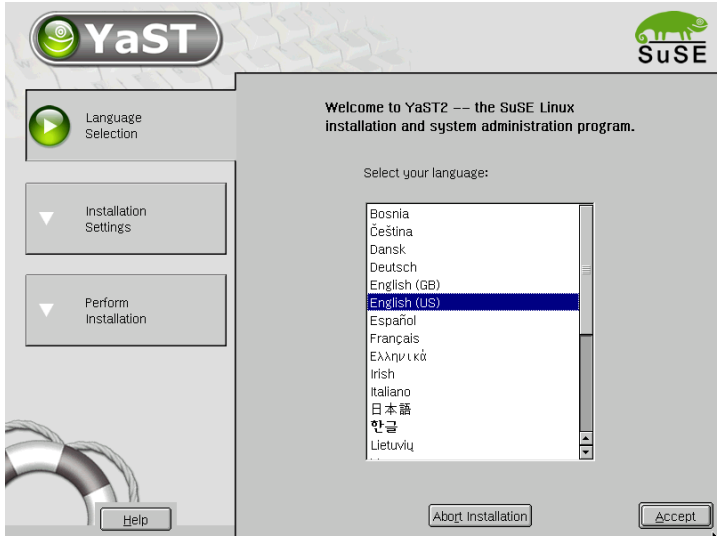

Figure 4.1: *Selecting the Language*

Installation Suggestions

After the hardware has been detected and the mouse configured, YaST2 provides information about the detected hardware and the suggestions for the installation and the partitioning. 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. You can also boot your installed system from here. This is useful if your system is no longer able to boot from the hard disk.

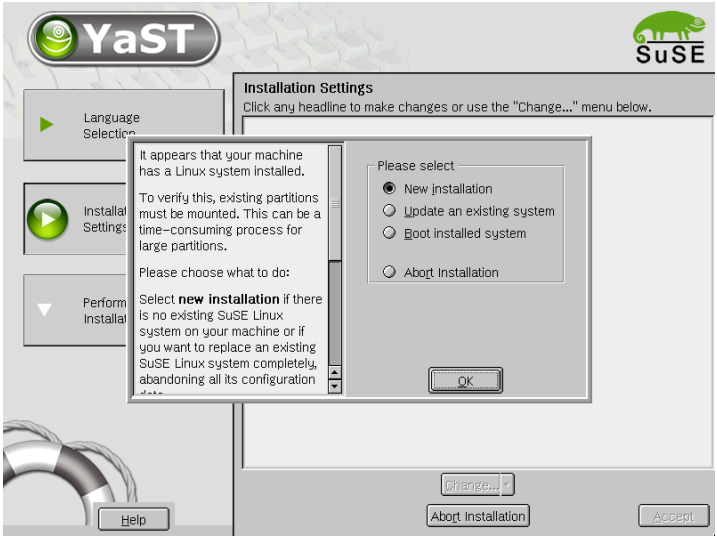

Figure 4.2: *Selecting the Installation Type*

Keyboard Layout (IBM pSeries only)

Select a keyboard layout. By default, the layout corresponds to the language selected. After changing the layout, use the field to test special characters, Ⓨ, and Ⓩ to make sure the layout is correct. If they are not displayed correctly, the keyboard layout is not correct. Return to the suggestions with 'Next'.

Mouse (IBM pSeries only)

If YaST2 did not recognize your mouse type automatically, use Tab until 'Change' is highlighted. Press Space and the arrow keys until 'Mouse' is selected. Press ↵ to open a mouse type selection screen, illustrated in Figure 4.3 on the following page.

To select your mouse type, use ↑ and ↓. Your mouse documentation should include a description of the mouse type. Select the mouse type from the list. Confirm your selection by pressing Alt + T or Tab and ↵.

Now, test to see if your mouse is working. If the mouse cursor on the screen follows your mouse movements, your mouse is configured correctly. If the cursor does not move, select a different mouse type and try again.

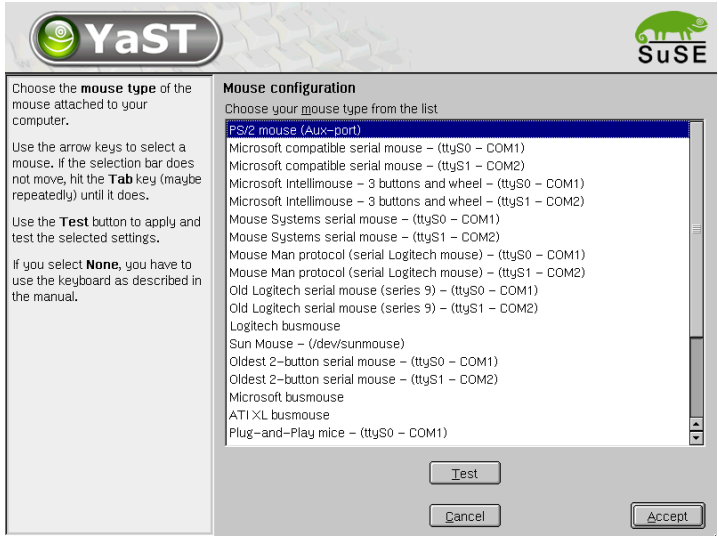

Figure 4.3: *Selecting the Mouse Type*

Partitioning

During installation you can split up the available disk space into several logical “partitions”. This procedure is called “partitioning”.

The YaST2 Partitioner

The suggested partitioning appears with a confirmation option, allowing you to accept, change, or discard the suggestion. See *More Partitioning Tips* on page 48 regarding automatic entries in the file system table. Clicking ‘Change’ opens the ‘Partitioner’ (Figure 4.4 on the facing page).

YaST2 shows all partitions present on the selected hard disk (Figure 4.4 on the next page). Any free storage space is shown and automatically selected. To assign more disk space to Linux, assign more partitions following the order from the last to the first partition. If you have three partitions, it is not possible to use the second one for Linux and the first and third for other operating systems.

If you choose ‘Discard’, a dialog appears containing the hard disk selection as shown in Figure 4.5 on page 46. All the hard disks on your system are listed

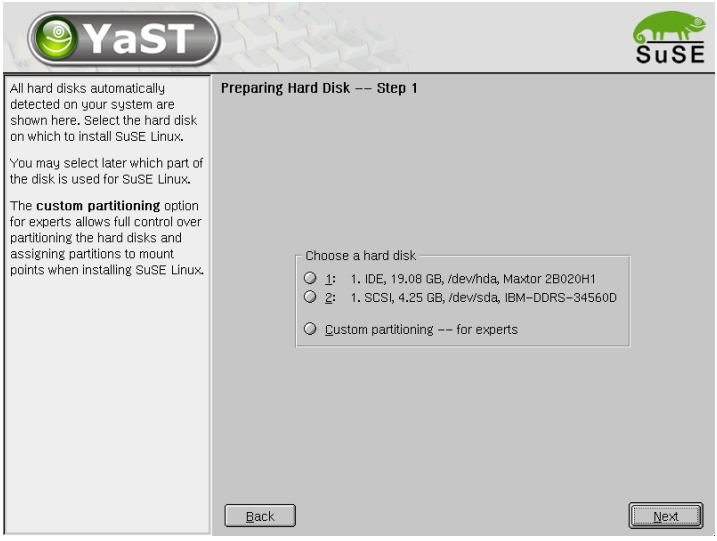

Figure 4.4: *Selecting the Hard Disk*

here. Decide where to install SuSE Linux Enterprise Server. There is no risk associated with selecting the hard disk, as no changes are made at this time.

'Advanced Settings, Manual Partitioning' shows the previous partitioning of your system. It can be manually modified or reused for your new system.

By clicking 'Entire hard disk' in the following dialog window, the entire hard disk and all its partitions are activated for Linux.

During the course of the installation, YaST2 will verify whether there is sufficient space for a minimum installation and whether the three standard partitions used by Linux can be created. If this is not the case, you will prompted to alter your selection. If there is enough storage space, YaST2 will use your settings to apportion the entire hard disk or the selected partitions.

Caution

If you choose 'Use entire hard disk', all information on your hard disk will be lost.

Caution

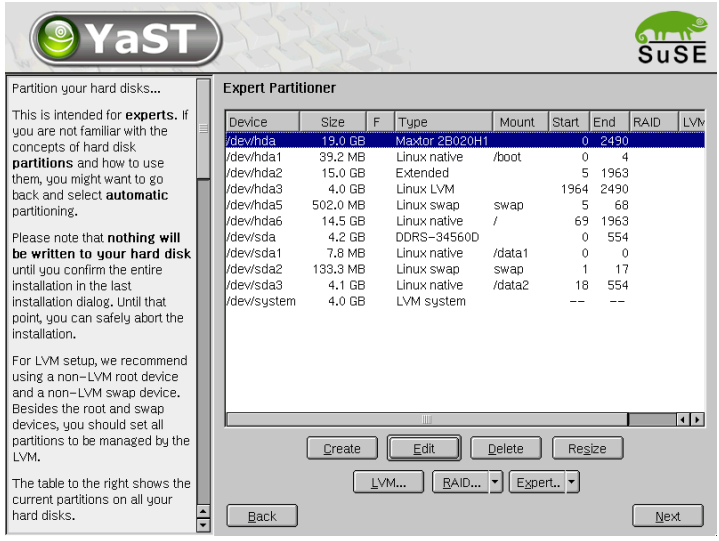

Figure 4.5: *The SuSE Linux Enterprise Server Partitioner*

Manual Partitioning

With the 'Partitioner', shown in Figure 4.4 on the page before, the partitions of your hard disk can be modified manually. Partitions can be added, removed, or changed.

If you select 'Partitioning' in the suggestion screen and 'Base partition setup on this proposal' in the next dialog, the partitioner will list the hard disk and all available or suggested partitions. Disks are listed as devices without numbers (such as `/dev/sda`). Partitions are listed as parts of the devices (for example, `/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. You can create up to four primary partitions or up to three primary partitions and one extended partition. Within the extended partition, you can create several "logical" partitions.

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 will be 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 do not want to format the partition): This can be 'Linux swap', 'Linux', 'Linux LVM', or 'Linux RAID'. For details on LVM and RAID, refer to the appendix, Chapter *Configuration with YaST2* on page 71.
   - 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 will also be 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/sdb1       /data1  auto    noauto,user 0 0
/dev/sdb8       /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.

Furthermore, you can select the following Softwarepacks:

'Minimal' This selection installs a fully functional Linux operating system providing textmode only and requires about 360 Megabytes of disk space. Selecting this type of installation does not install any of the typical "server" applications, like samba and ftp.

'Minimum graphical system (without KDE)' This option installs a minimal system including the X-Window-System. You will be able to use FVWM2 or Windowmaker as the windowmanager for the graphical environment. This installation requires about 500 Megabytes of disk space.

'Default System' This is the SuSE Linux Enterprise Server standard installation. This requires about 1.8 Gigabytes of disk space.

Select 'Details' to do a more detailed software selection.

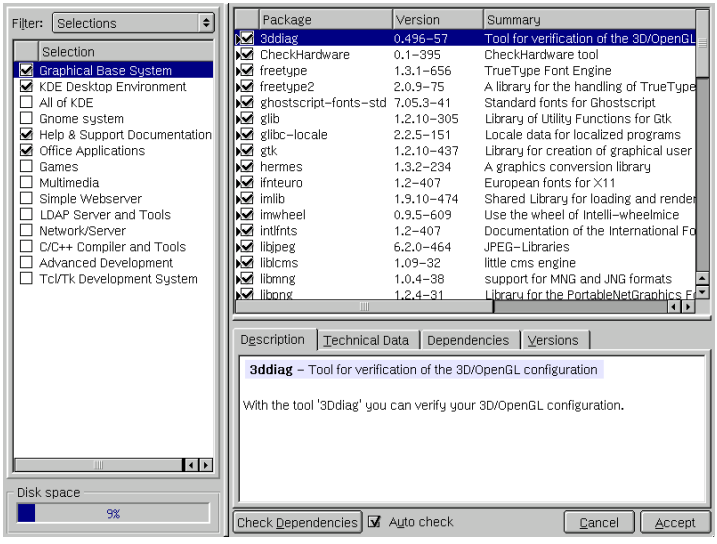

Figure 4.6: *Installing and Removing Software*

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. 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 'package groups', which is also defined as the default filter when you start the software selection in YaST 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.

Booting (IBM iSeries only)

The booting menu allows iSeries users to specify what locations they wish to have a kernel installed to and to customize the kernel parameters appended to the installed kernels.

By default all of the options available for your configured disk partitioning scheme will be selected. This means that you will, most likely, have the following kernels written to the following locations.

First, a stream file with the name listed in the 'iSeries streamfile for *STMF booting' field, will be written to your `/boot` directory in the installed system. This kernel could be transferred via FTP to your OS/400's IFS for '*STMF' booting in the future.

A kernel will also be written to the 0x41 PReP boot partition listed in the 'Choose 41 PReP boot partition' field. (NOTE: This option and field will not be available for users installing to native DASD only or for users that didn't create a PReP boot partition on their virtual DASD.)

Finally, on the OS/400 side a rescue kernel will be installed to slot A and a kernel for normal system usage will be installed to slot B.

Any combination of these options may be selected by checking or unchecking the radio boxes associated with each kernel installation option.

Any kernel parameters that you require to be added to all of these kernels can be added in the 'Kernel boot parameters' field.

If you do not want any kernels installed during installation you can choose 'Do not configure bootloader' and no kernel installation will be done by YaST2.

It is recommended that the user leave these settings at the default options provided, unless you are familiar with iSeries Linux and with how to compile and install updated kernels after installation has completed.

For information on how to use these kernels to boot your SuSE Linux Enterprise Server after installation see *Booting Linux* on page 59.

Time Settings

In this screen (Figure 4.7), choose between `Local Time` and `GMT` in the field marked '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.

Note

On IBM iSeries systems, the clock is automatically set to GMT as it is neccesary to interact properly with the LPAR's clock

Note

Starting the Installation

A click on 'Next' accepts the suggestions and any changes made. A green confirmation screen opens. After clicking 'Yes' here, the installation begins using your settings. The installation usually takes between fifteen and thirty minutes, depending on your machine's performance and your software selection.

System Configuration

After the installation of your system and selected software is complete, you need to make three more important settings before you can work with SuSE Linux Enterprise Server: define a password for the system administrator `root`, create a normal user, and configure your monitor. The following sections show how this is done.

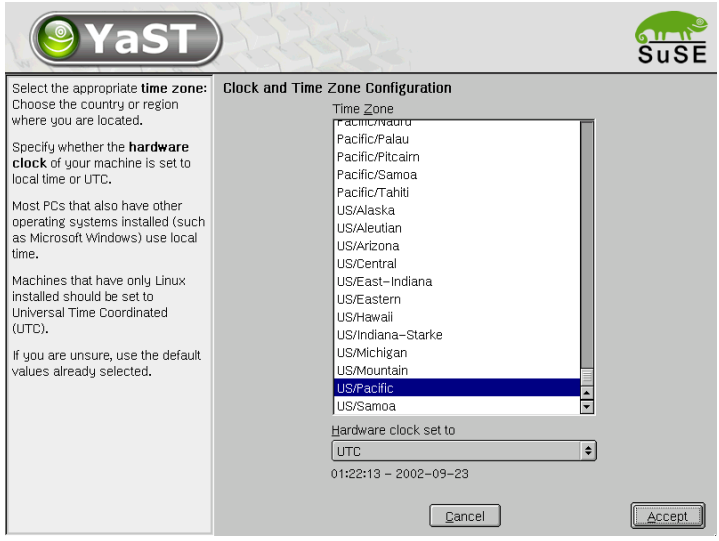

Figure 4.7: *Selecting the Time Zone*

Root Password

`root` is the name of the superuser, the system administrator. `root` is permitted to do all the things normal users are not permitted to do. The superuser can make changes to the system, such as installing new applications or setting up new hardware. If users forget their passwords or have problems with software, `root` can help them. 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 irreversibly any file on the system.

For verification purposes, the password must be entered twice as in Figure 4.8 on the facing page. Be particularly careful not to forget the `root` password. It cannot be retrieved later.

Caution

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 4.8: *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. For it to function smoothly, each user needs a user account, which allows them to log in to the system then to log out again when finished. Setting up user accounts provides a strong basis for operating security. Normal users cannot change or delete files needed for the system to work properly. Similarly, a user's personal data may not be accessed, modified, or deleted by other users. Each user can set up his own working environment with preferred applications and settings. When users log in to the Linux system, they find their personal environments unchanged.

Create such a user account yourself using the dialog as shown in Figure 4.9 on the next page. 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' and the system will automatically generate one for you.

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.

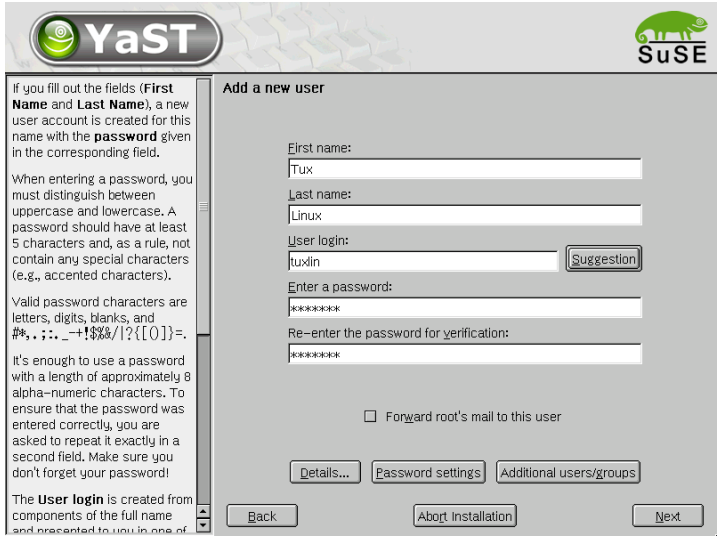

Figure 4.9: *Entering the User Name and Password*

Caution

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 8 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 the digits 0–9 may be used.

Caution

Monitor Settings (IBM pSeries only)

This shows graphics card and screen with a suitable configuration. In most cases, you can accept the suggestion. However, you can customize color depth, resolution, and the image repetition rate manually.

The settings will be tested once you have accepted the suggestion or entered your changes.

When you click ‘Change’, you have the option of configuring the graphical interface. For this purpose, the program SaX2 is started.

Hardware Configuration

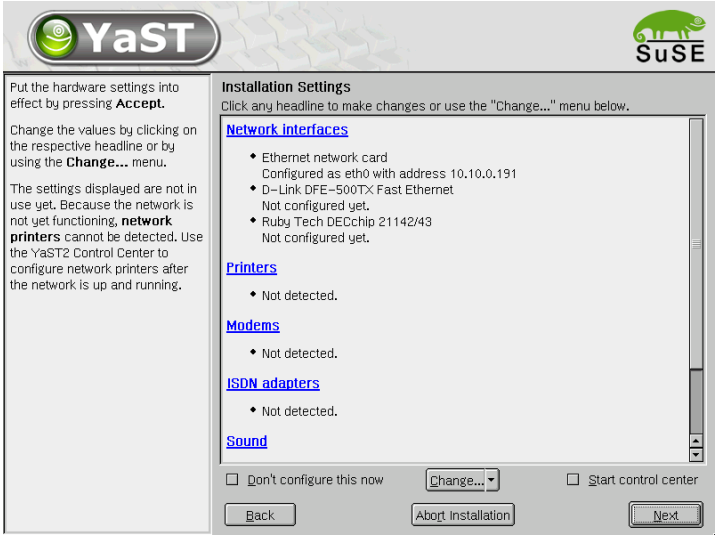

Figure 4.10: *Configuring the System Components*

Once your graphics card has been configured, you will see a screen like that shown in Figure 4.10. Now you have the option of configuring your system hardware, such as printer or sound card. The hardware configuration can also be done after the installation.

Start the hardware configuration by clicking each component. YaST2 will then automatically detect and configure the hardware.

When setting up network devices on iSeries, YaST2 will display any virtual ethernet devices as ‘IBM virtual ethernet card’. Click ‘Finish installation’ when completed.

If you have a mix of ethernet devices on your system, you must be aware of the device numbering. Write down which adapters are on the system, and how many of each. When a module is loaded, all the devices of that type of assigned ‘ethx’ device names. So the first device you configure as

'eth0', the other devices of the same type will take the 'eth1', 'eth2', etc, device numbers. Take the example of two IBM 79c970 ethernet cards (Feature Code 2838's using the 'pcnet32' driver) ethernet cards and two virtual ethernets on an iSeries. If you configure one of the IBM 79c970 cards as 'eth0' (this will happen automatically in this case), the other card will become 'eth1' when the system is booted. You should configure 'eth2' as your first virtual ethernet card, selecting the correct device number. Obviously, this can also happen if you have two different types of physical ethernet cards. The key is just to know your system configuration, and remember that once a module is loaded, all devices of that type are assigned device numbers. Then just use the network configuration tool to select the correct numbers, you do not need to configure every card, just make sure that a device number is reserved for it correctly.

Graphical Login (IBM pSeries only)

SuSE Linux Enterprise Server is now installed and configured so you can log in to your system. Your monitor displays the graphical login, as in Figure 4.11 on the next page. Enter the user name specified earlier and the respective password to log in to your system.

Where Next

Proceed with Chapter *Booting Linux* on page 59 to check and modify the boot configuration.

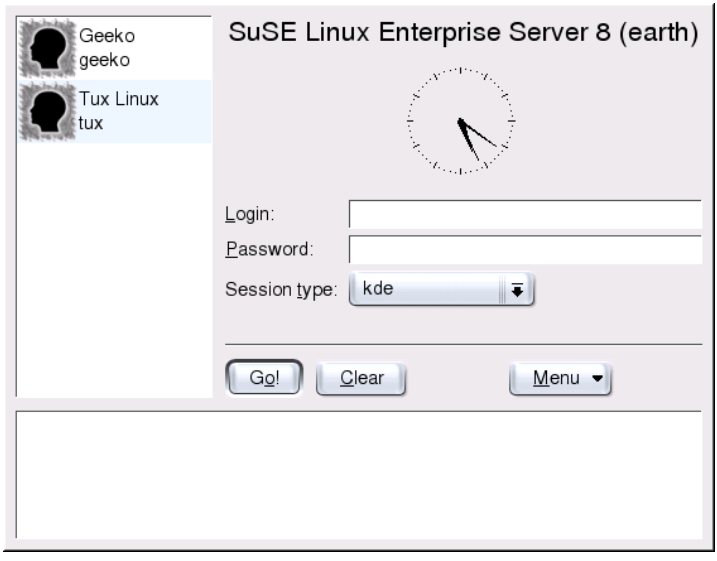

Figure 4.11: *Logging In*

Booting Linux

This chapter describes the procedure for booting an installed SuSE Linux Enterprise Server Enterprise Server for *iSeries and pSeries*.

Booting Linux on an iSeries System

The recommended method to power off an iSeries Linux partition is to simply vary off the NWSD. To reboot it, vary it back on.

After the installation, the NWSD must be configured to boot the installed system. There are three different boot methods:

Streamfile from IFS

As when booting the installation system, the IPL source can be a stream file — this time the plain kernel:

```
IPL source . . . . . .  *STMF__   *SAME, *NWSSTG, *PANEL...
IPL stream file  . . . .  '/KERNELS/ISERIES64'_________________
IPL parameters . . . . .  'root=/dev/hda3 noinitrd'___________
```

To boot a STMF kernel, you need to transfer the kernel from the Linux system, `/boot/vmlinux64`, to OS/400. You can ftp the kernel across the virtual ethernet or otherwise transfer it across the network. You can also use the install kernel, ISERIES64, but additionally give it the "noinitrd" IPL parameter option. It's best to copy the ISERIES64 file from the CDROM to IFS if you're going to do this.

Streamfile from Virtual Disk

Or, you can boot the kernel installed to the PReP partition on the first NWSSTG (virtual disk).

```
IPL source . . . . . .  *NWSSTG   *SAME, *NWSSTG, *PANEL...
IPL stream file  . . . .  *NONE_____________________________
IPL parameters . . . . .  'root=/dev/hda3'_________________
```

Kernel Slots

Otherwise boot from one of the two kernel slots (B in this case):

```
IPL source . . . . . .  B_____   *SAME, *NWSSTG, *PANEL...
IPL stream file  . . . .  *NONE_____________________________
IPL parameters . . . . .  *NONE_____________________________
```

If you boot from the A or B slot, the kernel command line is read from the LPAR 'Work with Partition Configuration' screen (option 14) in SST. You can also access and modify this through the proc file system on the Linux partition, by looking in `/proc/iSeries/mf/A|B/cmdline` (use the echo command like: `echo 'root=/dev/hda3 single' > cmdline` to modify it). For the A and B slots, the value in IPL Parameters is ignored.

Choosing a Boot Method One advantage of IPLing from a file in the IFS filesystem is that several NWSDs can boot an identical kernel. Native I/O only systems can be booted without an NWSD and are therefore less dependent on a fully running OS/400. If you get a different kernel from somewhere, such as a different, updated, linux partition, remember that you must install the modules on the partition for the system to work correctly.

This is discussed in depth in the IBM Redbook "Linux on iSeries" (SG24–6232–00)
(`http://www.redbooks.ibm.com/redpieces/pdfs/sg246232.pdf`).

- A NWSD with a linked NWSSTG can be booted from a PReP boot partition (IPL Source *NWSSTG), from one of the two kernel slots (IPL Source A or B), or from a stream file (IPL Source *STMF).
- A NWSD without a linked NWSSTG but with an attached native disk can be booted from one of the two kernel slots (IPL Source A or B) or from a stream file (IPL Source *STMF). A LPAR with attached native disk can be powered on from the A or B side in SST 'Work with Partition Status'.

If you have created a PReP boot partition, YaST has activated that partition and installed a kernel on it.

Whether you havea PReP boot partition or not, YaST2 will install a kernel to both the A and B kernel slots. YaST2 will install the ISERIES64 install kernel to the A slot (see *Booting (IBM iSeries only)* on page 50. The standard kernel (found in `/boot/vmlinux64`) will be installed to the B slot, with an attached ramdisk containing the ibmsis driver, if necessary. The install kernel in the A slot is designed to be a fallback if booting from the B side fails.

Some recommendations: If you just want to get started and do not know what you will want to use later on, consider these suggestions, in order:

1. If you have a PReP boot partition, select that.
2. If you have no virtual disk, select B.
3. If you have a strategy calling for everyone to boot from the OS/400 file system, select the *STMF option and point to `vmlinux64`.

Kernel Arguments: In the first and third case, we need to point the kernel to the root partition using the IPL Parameters field (`root=/dev/hda3` in this case). In the second case, this is achieved by the A and B command line slots read from the LPAR 'Work with Partition Configuration' screen (option 14) in SST. You can also access and modify this through the proc file system on the Linux partition, by looking in `/proc/iSeries/mf/A|B/cmdline` (use the echo command like such: `echo 'root=/dev/hda3 single' > cmdline` to modify it). For the A and B slots, the value in IPL Parameters is ignored.

By default, the kernel arguments in slot A are set to start in a recovery mode. The A slot command line looks like this after the install: start_shell manual=1 single. When a system is booted into single user mode, it can be brought to runlevel 5 with the command `init 3` (only root is allowed to log into the system in single user mode or use the init command).

First Boot: When booting the installed system, you might see some minor error messages, which you ignore for now. Also, depending on what was installed, YaST will start automatically upon first boot and run some scripts that configure the system. It will prompt you to press Enter a few times. Therefore, look at what is happening on the Console at first boot.

What Then When the system is up and running, you can install and update additional packages, like the kernel sources or other packages from the `updates` directory on the first CD-ROM or from our FTP site. See the next section for more information.

Booting Linux on a pSeries System

Thare are two ways to boot SuSE Linux Enterprise Server on a pSeries system:

- Using the zImage on the PReP partition
- Using yaboot

Using the zImage on the PReP partition

To use the PReP partition on your pSeries system to boot the Linux kernel image directly, copy the kernel image to this partition first.

```
earth:/ #  dd if=/boot/vmlinuz of=/dev/sda1 bs=4096
```

The PReP partition is `/dev/sda1` in this example.

Then check the boot configuration of your system's firmware. Enable boot from this PReP partition in the firmware. Consult the documentation provided with the hardware for details.

Using yaboot

yaboot gives greater flexibility to the boot process. With yaboot, you can

- configure several kernel images for booting
- select one of the configured kernels during the boot
- add an append line to the kernel for additional parameters

yaboot uses the configuration file `/etc/yaboot.conf`. This file is generated from the configuration file `/etc/lilo.conf`, which is similar to the one used on x86 systems. If you are already familiar with this configuration file, things will be much easier for you.

The LILO configuration file is `/etc/lilo.conf`. An example is File 2 (root here is `/dev/sda3` and the boot loader is on `/dev/sda1`).

```
boot=/dev/sda1
default=linux
timeout=100

image = /boot/vmlinuz
        root = /dev/sda3
        label = linux
        append = ""     # optional kernel parameters

image = /boot/vmlinuz.new
        root = /dev/sda3
        label = newlinux
        append = ""     # optional kernel parameters
```

File 2: *Configuration File* `/etc/lilo.conf`

Now modify the statements in `/etc/lilo.conf` to correspond with the data on your system. You can use `vi` to edit these. Specify the correct partitions and partition numbers. Use the entries as follows:

- at boot=/dev/sda1 enter your boot partition.
- default=linux specifies the label the boot loader uses when you press ↵ or after the waiting period (timeout) has expired.
- image = specifies the path to the kernel, assuming the root partition is correctly set.
- root = specifies the path to the root partition.
- append = "" transfers optional kernel parameters. Normally no entries are required here.
- image = adds another optional kernel to your configuration. This gives the ability to select a kernel during boot.

After adjusting the file `/etc/lilo.conf` to your system, start the following configuration by running LILO with `lilo`. After this, the boot configuration should be complete.

Updating the Kernel

When updated kernel rpms are distributed, you will need to think about how those upgrades affect how you are currently booting. In general, a kernel upgrade involves two steps: installing the kernel rpm, and then installing the kernel to your boot source (PReP, B slot, or STMF). If you are booting off of a PReP partition or the B slot, you shouldn't need any additional actions other than following the upgrade instructions (see the README with the upgrade for instructions).

For a STMF, you will need to install the kernel rpm update on the partition before booting the updated kernel. This ensures the modules get updated before booting the kernel. You must always first install the appropriate modules before booting a STMF kernel. If you don't, devices such as Ethernet cards where the driver is a module may not work correctly on the new kernel. Note that updating the kernel may affect your ability to boot the A slot recovery mode in any manner other than the recovery mode, but you can always boot the updated ISERIES64 with the recovery command line in the IPL Parameters.

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 A.1.

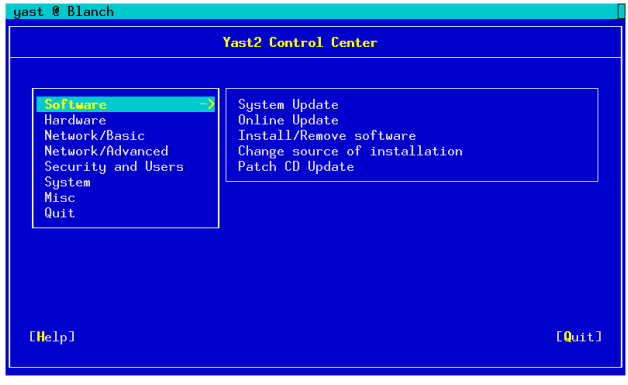

Figure A.1: *Main Window of the YaST2 Control Center*

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. A.2).

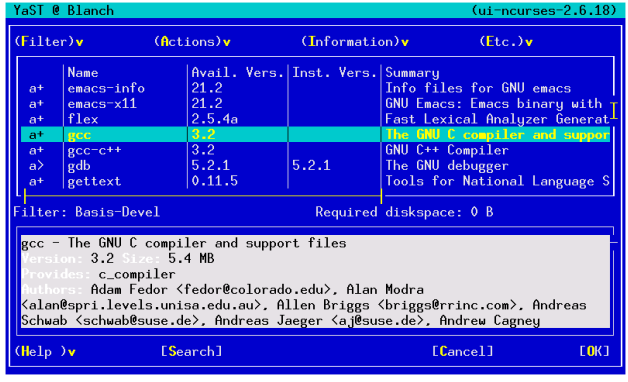

Figure A.2: *The Software Installation Module*

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` and then copy all the data.

The handling of LVM is most simple if you are going to set up a totally new logical volume. This can easily be done using YaST. After building a file system on the new volume, you can use it like any other partition.

If you need to add or remove space to a logical volume containing a running file system, you have to 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

Let´s have a look at the terms used by LVM. Knowing these terms should help to understand the different YaST menus.

Physical Volume (PV)

A PV is a physical medium (e.g. `/dev/sda`), which is prepared to be used by LVM. Therefore, some administrative data is added to it. To use a partition with LVM, the partition-type needs to be of the type “8E”.

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. You don´t have to worry about the PEs at this point, since 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 makes no difference 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 a LV is limited by the PE-size you 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, you need to choose another PE-size. For example, LVs containing 1 Terabyte can be achieved by using a PE-size of 16 MB.

How do I Access the LVs?

Assuming you have the 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 e.g. `test1`, `test2` ...`testn` you will find a hierarchy like

`/dev/suse/test1 /dev/suse/test2`

.
.
.

`/dev/suse/testn`

You can mount these volumes like regular partitions.

Further documentation:

Assumed, you have a running SuSE Linux Enterprise Server, you can 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

You can start the LVM configuration by selecting 'LVM...' in the 'Expert Partitioner' dialog.

The Partitioner

First, you will reach a dialog where you can change the partitioning of your hard disk. Here, remove or change current partitions as well as create new ones. A partition to use for LVM must have the partition label `8E`. These partitions are marked with the text "Linux LVM" in the partition list inside the window.

It is not necessary to individually set each partition designated for LVM to the partition label `8E`. When needed, YaST2 will automatically set the label of a partition assigned to an LVM volume group to `8E`. If there are unpartitioned areas on your disks, add LVM parititions in this dialog for all these areas. These partitions should immediately be set to the partition label `8E`. They do not need to be formatted and cannot be listed as a mount point.

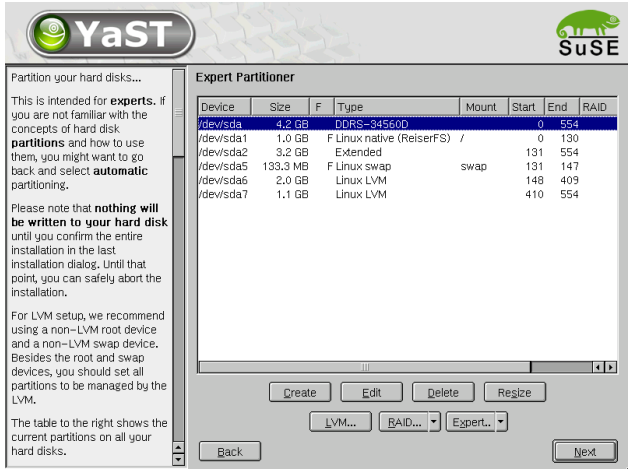

Figure B.1: *YaST2: LVM Partitioner*

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 being used. Of course, repartitioning the disks not attributed 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. On pSeries systems you need to use a root partition outside the LVM, unless you boot the system from a separate PReP partition. Select this root partition from the list and define it as the root file system using the 'Edit' button.

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

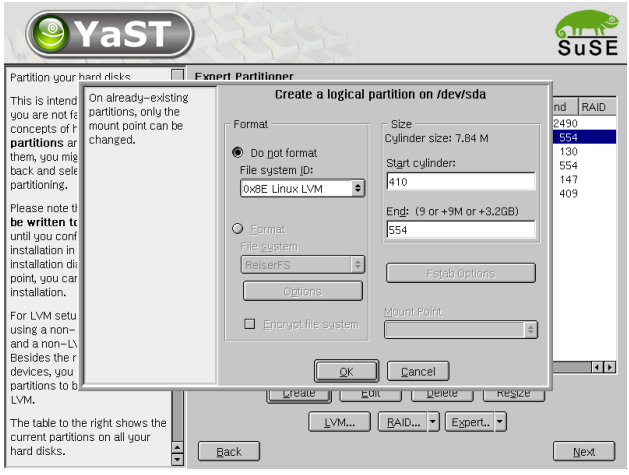

Figure B.2: *YaST2: Creating an LVM partition*

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 4 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.

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 will be 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 a previously unassigned partition to the volume group you selected, first select the partition and then click on the button ‘Add volume’ below the selection list. This allows the name of the volume group to be entered next to the

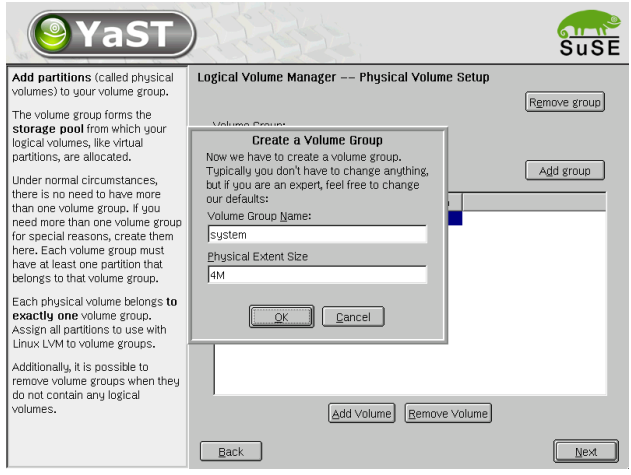

Figure B.3: *YaST2: Creating a Volume Group*

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 have 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 (with the 'Add' button), 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

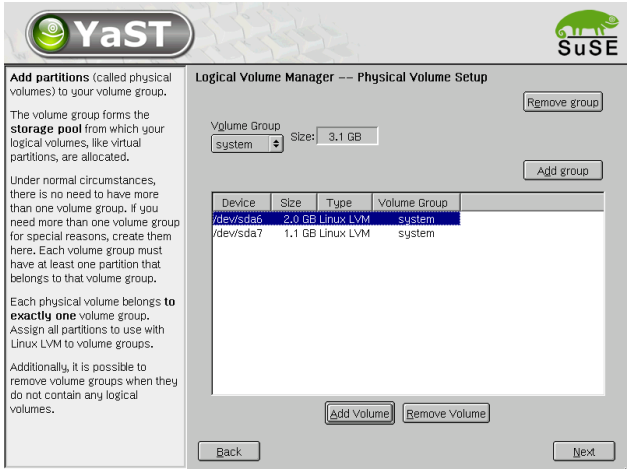

Figure B.4: *YaST2: Overview of Partitions*

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. You can then exit the dialog and continue on to software selection if you are currently in the process of installation.

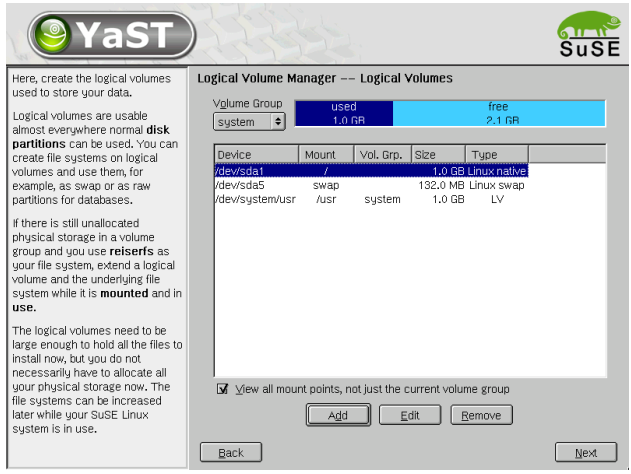

Figure B.5: *YaST2: Management of the Logical Volumes*

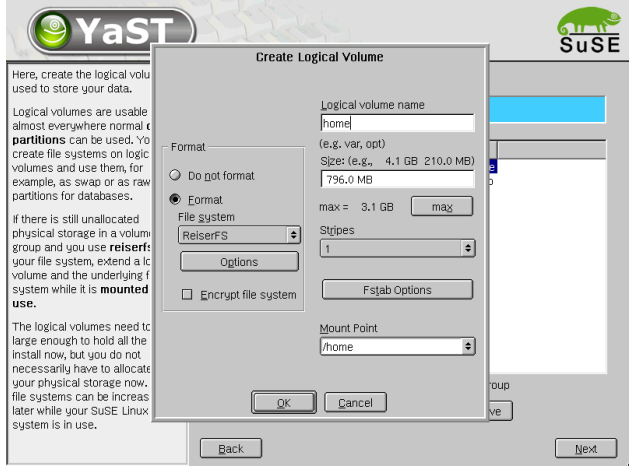

Figure B.6: *YaST2: Creating Logical Volumes*

File Systems in Linux

There are a number of file systems supported by Linux. 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 file system could be corrupted.

inode Inodes contain all sorts of 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 is actually stored.

journal In the context of a file system, a journal means an on-disk structure containing a kind of log where the file system driver enters what it is about to change in the file system's (meta)data. "Journaling" greatly reduces the recovery time of a Linux system as the file system driver does not need to initiate an exhaustive search for corrupt metadata all over the disk. Instead, the journal entries are replayed.

Major File Systems in Linux

Unlike two or three years ago, choosing a file system for your 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 in the world will never substitute for a reasonable backup strategy.

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. 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 (most of) the journaling file systems, e2fsck analyzes the whole file system and not just the recently modified bits of metadata. This takes significantly longer than the checking of 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. They can even be transformed while your Ext2 file systems are mounted. 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 has 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 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 in 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 in contrast to 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 for their AIX systems. 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 pro 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 two 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 C.1 on the next 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 originates from academic projects on operating systems and was the first file system used for Linux. Nowadays, it is used as a file system for floppy disks.
`msdos`	*fat*, the file system originally used DOS, used today 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 C.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 which are neither the owner nor members of the owning group). This scheme is sufficient in the majority of cases that occur practically; several workarounds to stretch its limits have been

used by system administrators for years. Advanced applications nevertheless occasionally require a more flexible scheme.

In such situations, when 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). Using 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 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/` or `http://www.suse.de/~agruen/acl/posix/posix.html`.

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 on which the file is stored. 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, and Ext3. 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 is contained 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 C.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

Table C.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 as long as no one tried to manipulate huge databases on Linux. Becoming more and more important for server computing, the kernel and GNU 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 C.3 on the next 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 Limit
Ext2 or Ext3 (1 kB block size)	16448 MB (~16 GB)	2048 GB (2 TB)
Ext2 or Ext3 (2 kB block size)	256 GB	8192 GB (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 C.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/ext2.html`
`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 follow 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. Maintenance 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

HB2112-3INST-INT