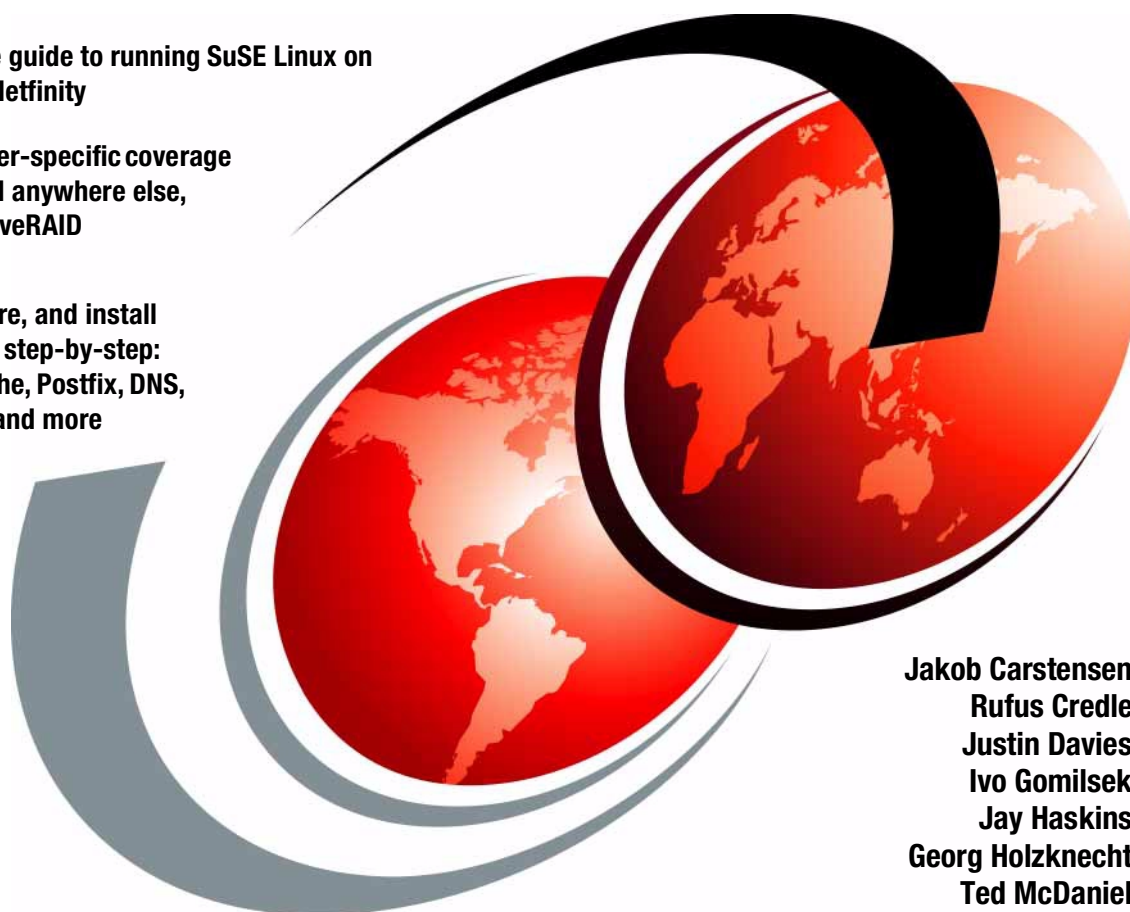


SuSE Linux Integration Guide for IBM **e**server[™] for xSeries and Netfinity

The complete guide to running SuSE Linux on
xSeries and Netfinity

Netfinity server-specific coverage
you can't find anywhere else,
including ServerAID
configuration

Plan, configure, and install
key services, step-by-step:
Samba, Apache, Postfix, DNS,
DHCP, LDAP and more



Jakob Carstensen
Rufus Credle
Justin Davies
Ivo Gomilsek
Jay Haskins
Georg Holzknacht
Ted McDaniel

ibm.com/redbooks

Redbooks



International Technical Support Organization

**SuSE Linux Integration Guide
for IBM @server for xSeries and Netfinity**

December 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, "Special notices" on page 385.

Second Edition (December 2000)

This edition applies to preparing for or installing SuSE Linux 7.0 Professional on xSeries and Netfinity systems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HQ7 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1999 2000. All rights reserved.

Note to U.S Government Users – Documentation related to restricted rights – Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	ix
The team that wrote this redbook	ix
Comments welcome	xi
Chapter 1. Introduction	1
1.1 The IBM commitment to Linux	1
1.2 SuSE Linux	2
1.3 Introducing the xSeries family of servers	2
Chapter 2. Installing SuSE Linux	5
2.1 Hardware considerations	5
2.2 Making the CD-ROM bootable	6
2.3 Linux installation with Yast1	7
2.3.1 Booting the installation system	9
2.3.2 Starting the installation	14
2.3.3 Partitioning and creation of file systems	16
2.3.4 Software package selection and installation	29
2.3.5 LILO - the Linux Loader	34
2.3.6 Time zone and clock settings	36
2.3.7 Network configuration	37
2.3.8 Mouse configuration	42
2.4 Installation using Yast2	46
2.4.1 Yast2: Booting the installation system	47
2.4.2 Yast2: Installation type	50
2.4.3 Yast2: Preparing the hard disk	51
2.4.4 Yast2: Software installation type	56
2.4.5 Yast2: LILO configuration	57
2.4.6 Yast2: Adding a normal user	59
2.4.7 Yast2: Committing your settings	61
2.4.8 Yast2: Installing the new system	63
2.4.9 Yast2: Configuring your display	66
2.4.10 Yast2: Configuring your devices	68
2.5 ServeRAID: Installation and configuration	69
2.5.1 Installing SuSE with the latest ServeRAID driver	70
2.5.2 Installing ipsutils.rpm	72
2.5.3 The ipssend commands	73
2.5.4 Replacing a defunct drive	86
2.5.5 Replacing a defunct drive with disabled Hot Spare Rebuild	87
2.5.6 Replacing a defunct drive with a hot spare drive installed	88
2.5.7 Using the ServeRAID Manager utility	94
2.5.8 Remote management of the ServeRAID adapter	98

2.6 XFree86 configuration	103
Chapter 3. Basic system administration	111
3.1 Adding and removing software packages using YaST	111
3.2 Package management using RPM	117
3.3 User and group administration using YaST.	118
3.4 Adding users on the command line.	122
3.4.1 Modifying users - the command line version	125
3.4.2 Deleting users - the command line version	125
3.4.3 Group administration using YaST	126
3.5 Network configuration with YaST	127
3.6 Changing the configuration file with YaST	132
3.7 System administration with Yast2	135
3.7.1 Yast2: Main window	136
3.7.2 Yast2: Network configuration	137
3.7.3 Yast2: NFS configuration	141
3.7.4 Yast2: Network services configuration	145
3.7.5 Yast2: Package maintenance	148
3.8 Finding Linux commands	149
3.8.1 File system permissions	150
Chapter 4. Samba	155
4.1 What can you do with Samba?	155
4.2 Setting up the Samba server	155
4.2.1 Configuring the Samba server	156
4.2.2 Starting and stopping the Samba server	164
4.2.3 Starting Samba as startup service	164
4.2.4 Using SWAT	164
4.3 Sources and additional information.	181
Chapter 5. DNS - Domain Name System	183
5.1 Installation of software	185
5.2 DNS sample configuration	185
5.3 Configuration tips	191
Chapter 6. Secure Shell	193
6.1 Installing SSH	193
6.2 Configuring SSH	193
6.2.1 Host key generation	194
6.2.2 User key generation	194
6.2.3 Configuring connections	195
Chapter 7. Apache and IBM HTTP Servers	197
7.1 The IBM HTTP Server	198

7.2 Apache HTTP Server installation	199
7.3 IBM HTTP Server installation	201
7.3.1 Activating IBM HTTPD on system bootup	203
7.3.2 Setting up the administration server	204
7.4 General performance tips	208
Chapter 8. Packet filtering with IP Chains	211
8.1 What is packet filtering?	211
8.2 What can you do with Linux packet filtering?	211
8.3 What do you need to run packet filtering?	212
8.4 Network configuration for a packet filtering implementation	212
8.5 How to permanently enable IP Forwarding	214
8.6 Your first IP Chains rule	214
8.7 How packets travel through a gateway	216
8.8 Using IP Chains	217
8.8.1 How to create a rule	218
8.8.2 Making the rules permanent	220
8.9 Sources of additional information	220
Chapter 9. DHCP - Dynamic Host Configuration Protocol	221
9.1 What is DHCP?	221
9.2 Why should I use DHCP?	221
9.3 Implementation on Linux	221
Chapter 10. Postfix	223
10.1 Installing Postfix	223
10.1.1 Postfix configuration files	224
10.2 Example Postfix installation	225
10.2.1 Telling Postfix about the server.	225
10.2.2 Telling the rest of the world about us	226
10.2.3 Transport and delivery options	227
10.2.4 Security goes a long way	228
10.2.5 Reporting Postfix errors	229
Chapter 11. NFS - Network File System	231
11.1 The NFS process	231
11.2 Allowing NFS access to data	233
11.3 Accessing data remotely with NFS - the command line view.	235
Chapter 12. NIS - Network Information System	237
12.1 What is NIS?	237
12.2 How can I use NIS?	237
12.3 Implementation on Linux	238
12.3.1 NIS server	239

12.3.2 NIS Client	242
12.4 Sources of additional information	245
Chapter 13. LDAP - Lightweight Directory Access Protocol	247
13.1 What is LDAP?	247
13.1.1 Directory Services	247
13.1.2 X.500	248
13.2 How can I use LDAP?	248
13.3 LDAP basics.	248
13.4 Implementation on Linux	249
13.4.1 Slapd.conf	250
13.4.2 ldap.conf.	250
13.4.3 nsswitch.conf	251
13.4.4 /etc/pam.d/login	251
13.4.5 Starting OpenLDAP	252
13.4.6 Testing authentication	253
13.4.7 Migrating /etc/passwd.	253
Chapter 14. General performance tools in Linux	257
14.1 General configuration hints.	257
14.1.1 Powertweak	258
14.1.2 Services	261
14.1.3 Kernel recompilation.	262
14.2 System monitoring and performance test tools	264
Chapter 15. Setting up a Beowulf cluster	275
Chapter 16. Backup and recovery	279
16.1 BRU	279
16.1.1 Installing BRU.	279
16.1.2 Basic commands	281
16.1.3 Basic backup	281
16.1.4 Basic restore	281
16.1.5 Basic verification and listing commands	282
16.1.6 X Interface	283
16.1.7 The big buttons in BRU.	283
16.1.8 Creating archives	284
16.1.9 Scheduling	286
16.1.10 Restoring files.	287
16.1.11 Listing and verifying archives	287
16.1.12 Summary	288
16.2 Microlite BackupEDGE	288
16.2.1 Installing Microlite BackupEDGE	289
16.2.2 Initializing the tape	290

16.2.3	Your first backup	292
16.2.4	Restoring single files or directories	296
16.2.5	Master and incremental backups	298
16.2.6	Restoring master and incremental backups	301
16.2.7	Performing scheduled backups	302
16.2.8	Configuring the tape devices	305
16.2.9	Defining the devices for making backups	311
16.2.10	Microlite RecoverEDGE	314
16.2.11	More information on Microlite	324
16.3	Arkeia	324
16.3.1	Installing Arkeia	325
16.3.2	Configuring Arkeia	325
16.3.3	Interactive backup	340
16.3.4	Periodic Backup	344
16.3.5	Restoration	345
16.3.6	Advanced features of Arkeia	350
Chapter 17. SNMP		351
17.1	Community Strings	351
17.2	Why should I use SNMP?	353
17.3	Implementation on Linux	354
17.3.1	MRTG	355
17.3.2	Sources of additional information	359
Appendix A. RAID levels		361
A.1	What is RAID?	361
A.1.1	RAID-0	362
A.1.2	RAID-1 and RAID-1E	363
A.1.3	RAID-10	364
A.1.4	RAID-5	365
A.1.5	RAID-5 enhanced	369
A.1.6	Orthogonal RAID-5	371
A.1.7	Performance	372
A.1.8	Recommendations	374
A.1.9	Summary	375
Appendix B. Sample smb.conf SAMBA configuration file		377
Appendix C. Special notices		385
Appendix D. Related publications		389
D.1	IBM Redbooks	389
D.2	IBM Redbooks collections	389
D.3	Other resources	389

D.4 Referenced Web sites	390
How to get IBM Redbooks	393
IBM Redbooks fax order form	394
Index	395
IBM Redbooks review	399

Preface

This redbook will help you install, tailor and configure the SuSE Linux 7.0 distribution on different servers of the xSeries and Netfinity class. You will be instructed on how to do the basic installation, and installing and configuring different services such as Apache (http-Server), Samba (Fileserver for Windows-based networks), and Postfix (an alternative to Sendmail), backup and recovery, and several other servers.

Linux is a very mature and stable operating system but the Linux kernel is constantly being updated in order to make the operating system better. This can make it difficult for Linux beginners, so be prepared for a bumpy ride and a steep learning curve. But it is worth the effort and, as they say at SuSE, don't forget to have a lot of fun...

The team that wrote this redbook



Figure 1. The team (left to right) Credle, Holzknrecht, Carstensen, Haskins, Gomilsek, Davies, (lower) McDaniel

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Jakob Carstensen is a Technical Support Marketing Specialist for Linux Solutions Marketing at IBM. He is also a former project leader of the International Technical Support Organization, Raleigh Center, where he managed residencies and produced redbooks. Before joining the ITSO, he worked in Denmark both for the IBM PC Institute teaching TechConnect and Service Training courses, and for IBM PSS performing level-2 support of Netfinity products. He has a Bachelor of Electronic Engineering degree and has worked for IBM for the past ten years.

Rufus Credle is a Senior I/T Specialist and certified Professional Server Specialist at the International Technical Support Organization, Raleigh Center. He conducts residencies and develops Redbooks about network operating systems, ERP solutions, voice technology, high availability and clustering solutions, IBM and OEM business applications, all running on IBM Netfinity and xSeries servers. Rufus's various positions during his IBM career have included assignments in administration and asset management, systems engineering, marketing and services. He holds a BS degree in Business Management from Saint Augustine's College. Rufus has been employed at IBM for 20 years.

Justin Davies is a systems administrator and product manager at SuSE UK. He has five years of Linux experience, and his expertise is in embedded Linux systems, systems administration and network integration. He joined SuSE in May of 2000 after graduating from the University of Derby, with a diploma in computer science.

Ivo Davies is an IT Specialist for Storage Area Networks and Storage in IBM Global Services - Slovenia for the CEE region. His areas of expertise include storage area networks (SAN), Storage, IBM Netfinity servers, network operating systems (OS/2, Linux, Windows NT), and Lotus Domino servers. He is an IBM Certified Professional Server Specialist, Red Hat Certified Engineer, OS/2 Warp Certified Engineer and Certified Vinca Co-StandbyServer for Windows NT Engineer. Ivo was a member of the team that wrote the redbook *Designing an IBM Storage Area Network*, SG24-5758, *Implementing Vinca Solutions on IBM Netfinity Servers*, SG24-5843 and the first edition of various Netfinity and Linux Integration Guides. He also provide Level 2 support for IBM Netfinity servers, high availability solutions for IBM Netfinity servers and Linux. Ivo has been employed at IBM for four years.

Jay Haskins is a Systems Architect for IBM Global Services Enterprise Architecture and Design in Seattle, Washington. He has been a Linux and Open

Source advocate for more than five years and currently spends most of his time developing dynamic monitoring tools using Perl and the Apache Web server. Before joining IBM, Jay worked in several different areas of the information technology field including UNIX system administration, database design and development, Windows application development, and network administration.

Georg Holzknrecht is a Senior System Consultant at DeTeCSM, Darmstadt, Germany. He has 30 years of experience in different areas of the information technology field. He holds a diploma degree in electrical engineering from Technische Hochschule, Darmstadt. His areas of expertise include system programming for mainframes, network operating systems (NetWare, Linux), database administration and design, application and driver development, and systems management solutions with Tivoli.

Ted McDaniel is a Senior Support Specialist at the IBM PC HelpCenter in Research Triangle Park, NC. He is the World Wide Level 2 Linux support leader for IBM x-Series and Netfinity servers. Ted has six years of experience with Level 2 support.

Thanks to the SuSE development team for their support and a great distribution.

Thanks to original authors, Lenz Grimmer and Joe Kaplenk, for their contribution to the first edition of this redbook, which was titled *SuSE LINUX and Netfinity Server Integration Guide*, SG24-5863-00.

Thanks to the following people for their invaluable contributions to this project:

Diane O'Shea, Gail Christensen, Linda Robinson, Margaret Ticknor, and Tamikia Barrow
International Technical Support Organization, Raleigh Center

Thanks to the following people for their support:

Ruediger Berlich, Jasmin Ul-Haque, Malcom Yates, Jens Axboe, Joseph 'Uzi' Uziel, Dave Jones, Rafiu Fakunle and Rachael Edwards.

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Please send us your comments about this or other Redbooks in one of the following ways:

- Fax the evaluation form found in “IBM Redbooks review” on page 399 to the fax number shown on the form.
- Use the online evaluation form found at ibm.com/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Introduction

Linux is a UNIX-like open-source operating system and was the original creation of Linus Torvalds from Helsinki, Finland in 1991. He wrote the first kernel, the underlying program interfacing and running the computer hardware. Torvalds invited programmers from around the world to comment on and to improve his code. This is one of the key ideas behind the success of Linux. With the world as your laboratory, the number of testers and developers is nearly endless. It is because of this resource that Linux is constantly evolving and improving.

With the Linux source code being freely available, several companies have developed different distributions of Linux. A distribution is a complete system. The key component is the Linux kernel. Other utilities, services, and various applications can be included as well, depending on the distribution and the intended use. There is no standard distribution. Each of the many distributions available has unique advantages.

IBM was early to recognize the value of Linux, investing in Linux-related product development, forming alliances with key Linux distributors, contributing to the open-source community, and aggressively supporting the platform. IBM believes this investment will benefit its customers as they continue to exploit Linux for their IT infrastructures and e-business applications.

1.1 The IBM commitment to Linux

IBM is fully committed to the open source movement and believes Linux will emerge as a key platform for e-business. IBM will work with the open source community, bringing relevant technologies and experience to the table to help enhance Linux, to define the standards and to extend Linux to the enterprise level. IBM provides continued support and participation in three main locations:

- The Open Source Development Lab
- IBM Development and Competency Centers for Linux
- IBM Technology Center

As part of this continuing commitment, IBM has teamed with leading commercial Linux distributors, Caldera Systems, Red Hat, SuSE, and TurboLinux to port, test, and certify the performance of IBM offerings running on various Linux distributions, enabling you to exploit the full potential of Linux.

1.2 SuSE Linux



Founded in 1992 by four German software engineers, SuSE (pronounced soo'-sah) is the oldest major Linux solutions provider. The company's commitment to the highest standards in open source software development has helped the award-winning SuSE Linux operating system become one of the most widely used Linux distributions in the world.

Today, SuSE has an international presence, with offices in Europe, Latin America and the United States. SuSE Inc., the company's North American subsidiary was established in Oakland, California in 1997 and provides solutions, technical support, Premier Partner programs, and sales support to customers, resellers and distributors.

SuSE's mission is two-fold: to make open source solutions the global standard and to establish itself as the most reliable source for Linux on a worldwide basis. SuSE is a chosen IBM distribution partner for Linux.

The IBM @server xSeries Brand team works closely with SuSE and other distribution partners to fully test and certify xSeries and Netfinity servers are ready to perform with Linux. Additionally, xSeries servers are designed with X-architecture technologies to bring reliability, high availability, powerful performance and manageability in an affordable Intel processor-based platform with products to meet a variety of customer needs and choices of operating environments.

1.3 Introducing the xSeries family of servers

IBM @server xSeries is the new IBM Intel server brand. xSeries are Intel processor-based servers with X-architecture technology enhancements, for a level of reliability, performance and manageability previously out of reach for industry-standard servers. This represents a full circle of technology evolution for Netfinity heritage in X-architecture, which is based on technologies derived from the IBM ES, RS and AS series servers, bringing mainframe category technology to the industry-standard architecture. Also, NUMA-Q will be aligned with xSeries to ensure IBM resources are focused most effectively on the Intel marketplace.

xSeries servers are available in the following four categories:

- Point Solution Servers
- Universal Servers
- Rack Optimized Servers
- Extremely Scalable Servers

For more information on the xSeries, visit the Web site at:

<http://www.pc.ibm.com/us/eserver/xseries/>

Chapter 2. Installing SuSE Linux

This chapter discusses the basic installation of SuSE Linux 7.0 Professional on different models of IBM Netfinity servers and how to work around common problems. Since it is almost impossible to cover all hardware combinations, we have concentrated on typical configurations, which are representative examples:

- IBM Netfinity 3000
- IBM Netfinity 3500 M10
- IBM Netfinity 5000 with ServeRAID controller
- IBM Netfinity 5500 with ServeRAID controller
- IBM Netfinity 5600 with ServeRAID controller
- IBM Netfinity 7000 with ServeRAID controller
- IBM Netfinity 8500 with ServeRAID controller

We strongly recommend that you also have a look at the extensive SuSE manual, which covers the installation process in more detail and more variations than we will describe here. It also gives you a lot of background information to begin with. Before you start the installation, make sure that you check the SuSE Web site for updates and bug fixes. Linux is a fast-moving target, and the development is a continuously ongoing process. There might be new boot floppy images or kernel patches that contain newer drivers. Also, make sure that you add all security fixes if you plan to connect your machine to the Internet. Updates and bug fixes for SuSE Linux 7.0 can be found at:

http://www.suse.de/en/support/download/updates/70_i386.html

The updates are located on the SuSE FTP server at the following address:

<ftp://ftp.suse.com/pub/suse/i386/update/7.0/>

2.1 Hardware considerations

Before installing SuSE Linux, it is helpful to know the hardware components in the computer that will be used for the installation. SuSE Linux is capable of detecting most of these components correctly. However, you should still try to gather information about the following components of your machine:

- **SCSI adapter** - manufacturer and model number
- **Hard drives** - interface type (SCSI or IDE) and size
- **CD-ROM** - interface type (SCSI or IDE)

- **Display Adapter** - manufacturer, model and video memory size
- **Mouse** - mouse type and connector type
- **Network card** - manufacturer and model
- **RAM** - the amount of random access memory in your system
- **CPU** - the type and number of processors
- **Monitor** - manufacturer and model, horizontal and vertical frequency range

A very helpful resource for information about IBM Netfinity servers and other IBM products including monitors and SCSI adapters can be found on the following site:

<ftp://ftp.pc.ibm.com/pcicrse/psref>

This archive contains Personal Systems Reference sheets (PSREF) for all IBM PC products, current and withdrawn. You can also get a lot of useful information about IBM hardware at the following Web sites:

<http://www.pc.ibm.com/support/>

<http://www.pc.ibm.com/us/eserver/xseries/library/index.html>

SuSE also maintains an online database of supported hardware for Linux, which is available at:

http://cdb.suse.de/cdb_english.html

In addition to that, SuSE certifies IBM Netfinity systems for compatibility with SuSE Linux and works closely with the developers at IBM.

2.2 Making the CD-ROM bootable

If you plan on booting the system directly from the CD-ROM, make sure the CD-ROM drive is the initial boot device prior to the installation. This can be accomplished by following these steps:

1. Power on the server.
2. When you see the IBM logo press F1 to enter the setup utility.
3. From the setup utility select **Start Options**.
4. From the Start Options select **Startup Sequence**.
5. Make sure that your CD-ROM is the initial boot device.
6. Press Esc until you see the setup utility main window and select **Save Settings**.

7. Press Enter to confirm saving the current settings.
8. Exit the setup utility.

Note

Making the CD-ROM bootable can also be done by loading the default settings from the setup utility, but be aware that all other settings will be set to default as well.

2.3 Linux installation with Yast1

SuSE Linux 7.0 Professional allows two types of installation. The first way is to use Yast1, the console-based installation. The other is to use Yast2, a GUI installer that allows a quick and easy way to install SuSE.

Yast2 is now commonly used to install SuSE and has been tested on many of the Netfinity range of servers. There are underlying problems with the S3 video chipset that in some situations makes installation with Yast2 infeasible. In these cases, using Yast1 for installation is the only option.

Note

If you have any trouble while installing on a Netfinity with ServeRAID installed, try booting the installer using:

```
ftp://ftp.suse.com/pub/suse/i386/update/7.0/kernel/ips-4.40
```

See section 2.7.3 in the SuSE manual for instructions on how to install the floppy boot image.

The ips-4.40 directory also contains the ips.o module for ServeRAID on uniprocessor systems, as well as ips-smp.o for multiprocessor systems.

For information on how to update the ServeRAID driver for the running system and install SuSE, see the ServeRAID section at the end of this chapter.

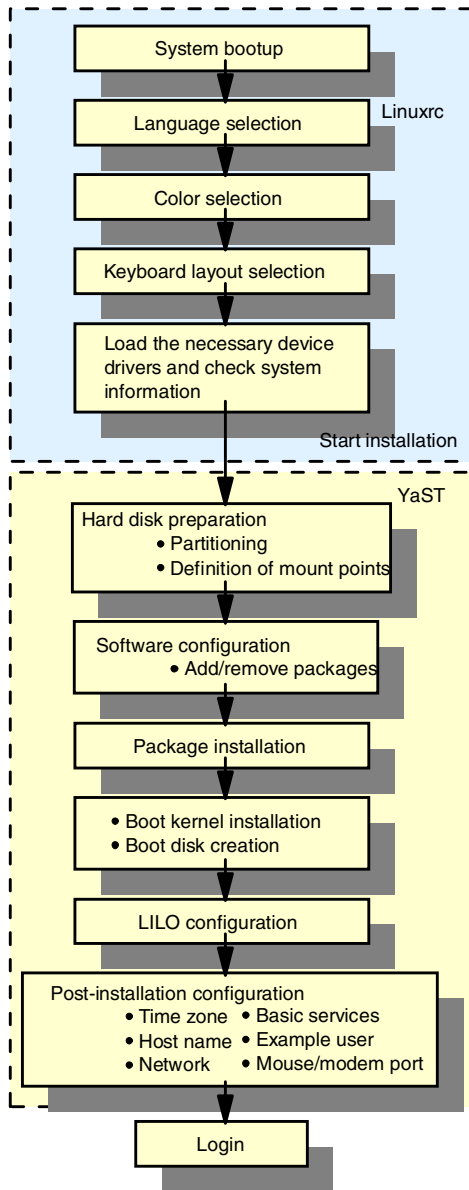


Figure 2. SuSE Linux installation workflow

The installation process is performed by two programs. Stage 1, the initial bootup process and the loading of driver modules, is fulfilled by Linuxrc, which can be loaded from either floppy disk or the CD-ROM. Linuxrc will then start YaST (Yet another Setup Tool) to perform such tasks as partitioning,

definition of mount points and installing the software packages. YaST can either be loaded from a local installation medium such as the CD-ROM or the hard disk, or it can be loaded over the network to do a network installation. After YaST has finished its job, it will boot directly into the freshly installed system. However, YaST will not only be used for the initial installation. You can use it for system administration and package management later on. Start it by typing “yast” as the root user. See Chapter 3, “Basic system administration” on page 111 for details.

2.3.1 Booting the installation system

Insert the second SuSE Linux 7.0 CD-ROM in your CD-ROM drive (use the first CD and the boot floppy disk, if booting from CD-ROM is not supported) and reboot or power up the system. After a short amount of time you should see the bootup splash window shown in Figure 3. If this fails, make sure that the system is correctly configured for booting from CD-ROM or floppy.

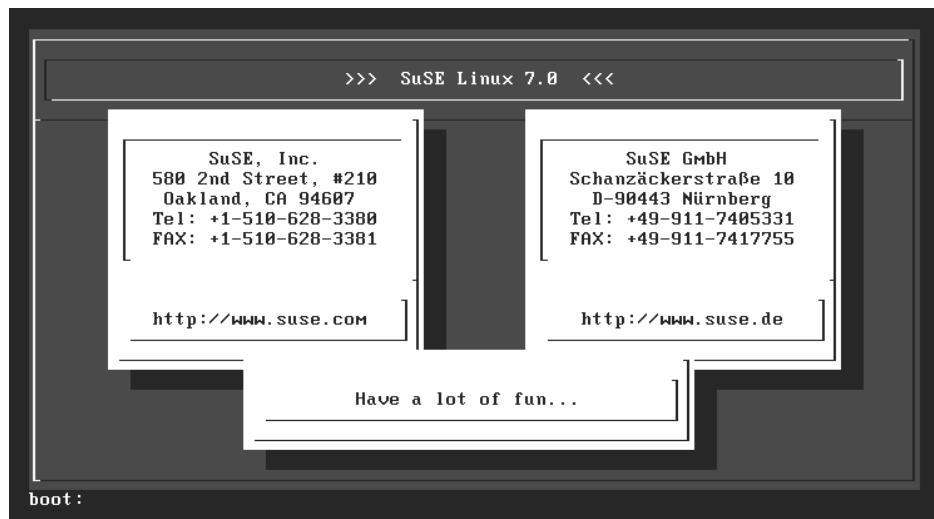


Figure 3. SuSE Linux bootup splash window

The boot prompt enables you to enter special boot parameters. This may be necessary if the system does not recognize certain hardware components. Section 14, “Kernel parameters” in the SuSE manual gives you more information about this feature. We did not experience any problems with IBM Netfinity hardware; therefore, you should not need this.

The installation system will automatically continue the boot process after a few seconds. If you press Enter, it will boot up immediately.

You will be asked for CD 1 after the system has booted. Insert it and click **Continue** to proceed with the installation.



Figure 4. Language selection window

First you have to select your desired language. Use the Up/Down cursor keys to highlight your selection and click **Ok** to continue.



Figure 5. Keyboard selection window

Now you have to select the required layout for your keyboard. Click **Ok** to advance to the Linuxrc main menu. It is very important that you select the correct keyboard layout. Choosing an incorrect keymap will result in unpredictable behavior from your keyboard, and may result in you not being able to continue with the installation.




Figure 6. Linuxrc - main menu

Figure 6 shows the Linuxrc main menu. It offers the following options:

Settings - This option enables you to modify the language, window or keyboard settings, if you need to revise the selection you made during the bootup process.

System information - This menu option gives you detailed information about the hardware that has already been recognized.

Kernel modules (Hardware drivers) - Use this menu to load device drivers for special SCSI devices, network cards and other devices.

Start installation / System - After you have loaded the necessary device drivers, select this option to continue the installation.

End / Reboot - This aborts the installation and reboots the system.

Before you can start the installation, you should make sure that the system detected your hard disk(s) and CD-ROM drive. If you intend to make a network installation or if you want to use a network connection later on, you should also load the respective network driver. Select **System Information > Hard disks / CD-ROMs** to determine which devices have been detected.

Devices that are connected to the Adaptec SCSI Host adapter, which is used in most IBM Netfinity servers, should already show up in this list. Return to the main menu and select **Kernel modules (hardware drivers)** to load the network and additional SCSI drivers.




Figure 7. Hardware driver selection window

Select **Load network card module** to load the network card driver and any other system drivers. You can also select **Autoload of modules** to let the system try to automatically probe for additional devices. However, this may freeze the machine or may not detect all components.

If your system has only a ServeRAID controller, it is imperative that you load the ServeRAID driver (ips.o). On the machines we tested, the ServeRAID drive was automatically loaded when we chose **Autoload of modules**. If your system locks up when autoprobing, you will need to manually load the module from the **Load SCSI module** menu item.

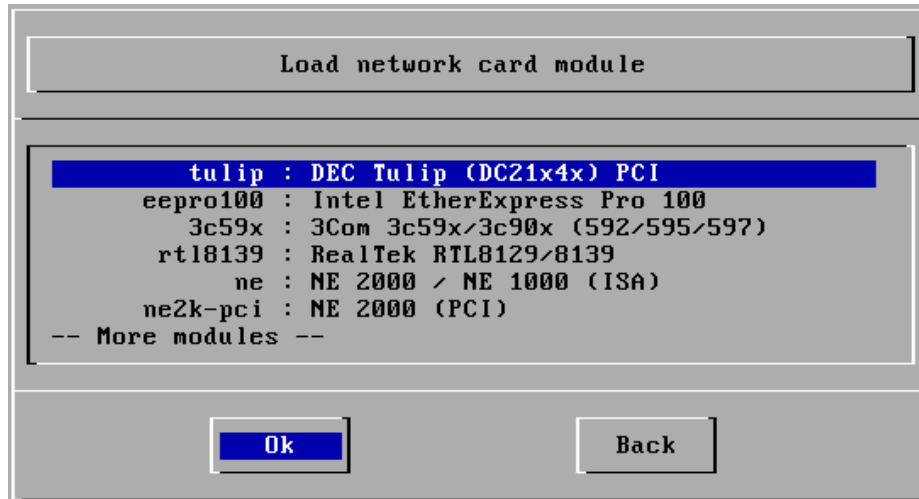


Figure 8. Network module selection window

Load the network card module that fits your network card. Select **eeepro100** if you have an Intel network card, or select the **pcnet32** driver from the separate modules disk if your PC uses a card with the AMD chipset. Before loading the driver, you can pass parameters to it (for example, interrupt and I/O address). This is not necessary for most modern PCI cards; you can just click **Ok** here. Linuxrc will now attempt to load the kernel module and will inform you of the success or failure including the output of the device driver startup. This procedure may take a while with some drivers, so you will have to be patient if the system does not react immediately.

After loading all necessary drivers, select **Back** to return to the Linuxrc main menu shown in Figure 6 on page 11. Select **Start installation / system** to begin the installation.

2.3.2 Starting the installation




Figure 9. Linuxrc: start installation

Your choices on this window are:

Start installation - to begin a regular installation.

Boot installed system - comes in handy if an already installed system fails to boot from the hard disk and you do not have a special boot disk.

Start rescue system - enables you to start a minimal Linux system in a RAM disk, which you can use to do system maintenance or repair a corrupted installation.

Start Live CD - enables you to run a full-fledged Linux system (including XFree86, KDE and compilers) directly from CD-ROM without installing Linux on your hard drive. You need to have the special Live-CD-ROM to do this, which is a separate product and is no longer included in the SuSE Linux box.

Eject CD - This will eject the CD in the system. It is good for those hard-to-reach eject buttons.




Figure 10. Selection of the installation medium

This window allows you to choose your source medium. In our case, select **CD-ROM**. YaST will load and continue the installation. You can also set up a file server that serves the installation CD-ROMs over the network using NFS or FTP. However, this is beyond the scope of this manual and will not be discussed. Please see the SuSE manual for further details about this.

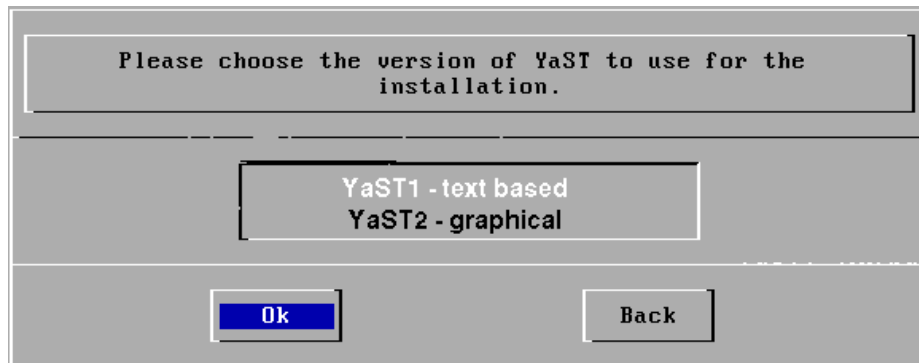


Figure 11. Choose Yast installation type

When you are asked which installation type you wish to proceed with, choose **Yast1 - text based**.



Figure 12. Type of selection

Select **Install Linux from scratch** to advance to the next section. If you intend to update an existing SuSE Linux installation, use **Update existing Linux system** here. Do not try to update distributions other than SuSE Linux with this feature! This can cause chaos in your installation. Choosing **Installation using Expert mode** gives you some more control over the installation process, but will not be discussed here.

2.3.3 Partitioning and creation of file systems

In order to be able to install Linux on your hard drive, you need to have some free space on your hard disk. This free space has to be divided among several partitions. Similar to fdisk in MS-DOS/Windows, SuSE Linux provides a tool to create the partitions and define their size and the partition type. After you have created the partitions, file systems have to be created on them (they need to be formatted) so that Linux can access them.

Linux does not know about drive letters such as A:, C: or D:. Everything lives below a single directory tree (the root directory). File systems on other partitions will be mounted to a subdirectory of the root directory. You will also have to define these mountpoints when creating the file systems on your partitions.

Devices also use a different naming scheme from the Microsoft operating systems. Instead of using drive letters, all drives in Linux are named alphabetically. Each partition on this drive has another number (CD-ROMs do not have partitions). For example:

- /dev/hda is the first Integrated Drive Electronics (IDE) drive (master on the first IDE channel).
- /dev/hdc would be the first IDE drive on the second IDE channel.
- /dev/hdb1 is the first primary partition on the slave drive of the first IDE channel.

- /dev/sda names the first Small Computer System Interface (SCSI) hard disk.
- /dev/sdb5 names the first logical partition on the second SCSI disk.

For more information about devices, see , “Device files in the /dev directory” in the SuSE manual.

Note

Even though the partitioning tool is capable of creating partitions for MS-DOS or Windows, you should not use it for creating partitions for operating systems other than Linux. Use the version fdisk that ships with MS-DOS/Windows to create such partitions.




Figure 13. Selection: Hard drive partitioning

At first you will be prompted if you want to create partitions on your hard disk(s). Select **Partitioning**, if you have not defined any partitions for Linux yet. If you want to keep previously defined partitions, choose **Do not partition**.

Setting up LVM allows you to configure the Logical Volume Management subsystem. LVM allows you to define a set of disks to appear as one unified disk system. More information on this can be found at:

<http://linux.msede.com/linux>



Figure 14. Selecting partitioning method

If you have multiple hard disks, you will be prompted to select the drive you want to use for the partitioning. You will then return to this window, after you have created partitions on one of the disks to be able to partition the other disks as well. Linux can be spread over multiple disk drives without problems.

If you choose **Whole hard disk** here, YaST will automatically partition the selected disk for you by creating one swap partition, a small partition for the /boot directory and one large partition for the root directory and all its subdirectories. It will also automatically define these mount points and advance to the package installation menu shown in Figure 25 on page 29.

While automatic partitioning is fine for home or workstation use, you should consider partitioning your disks manually to better fit your needs. To continue with manually partitioning your hard disks, select **Partitioning**.




Figure 15. Fdisk main window

If the current hard disk has not been used before, you will start with an empty partition table. You can now start adding partitions with the F5 key. Use F4 to delete previously defined partitions.

Note

Partitioning your hard disks is highly dependent on the purpose of your system. Depending on the intended services, you may need to create one especially large partition (for example for a file server). There is no general rule for this and it's almost impossible to give recommendations. See section 2.10, "Partitioning for experts" in the SuSE manual for more information on this issue.

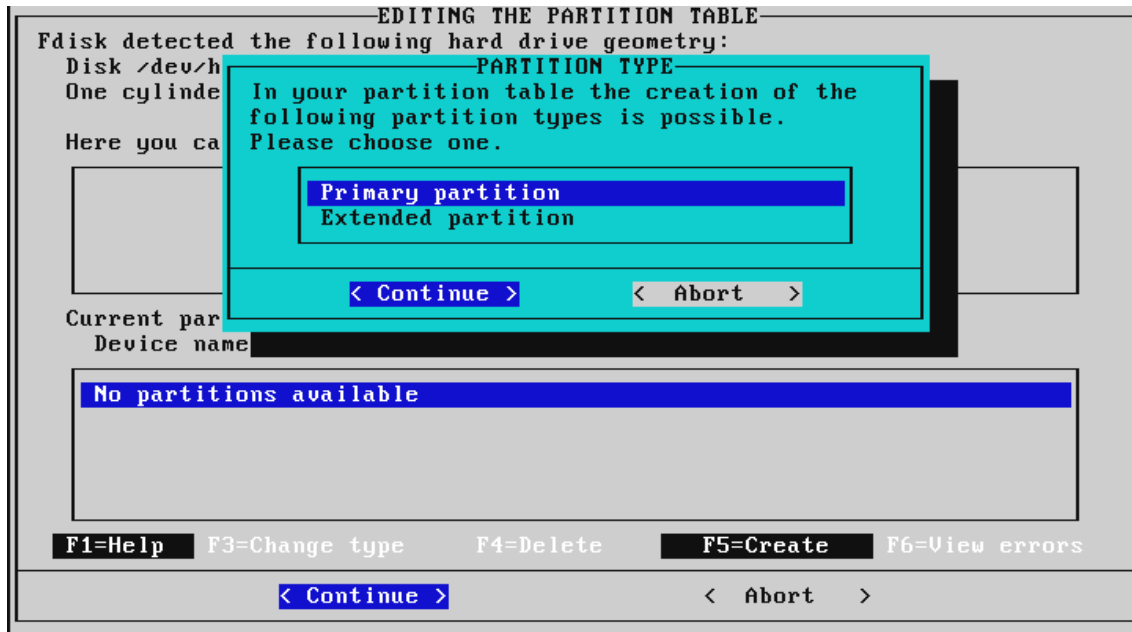


Figure 16. Selecting partition types

Depending on the already existing partitions, you can now define the partition type. A hard disk can consist of a maximum of four primary partitions, or up to three primary and one extended partition. An extended partition can contain multiple logical partitions. See section 3.3.9, “Partitioning your hard drive” in the SuSE manual for a detailed description of the different partition types on a PC. Linux can be installed in either partition type.




Figure 17. Configuring partition size

After defining the partition type, you now have to enter the size and physical location of that partition by supplying the starting and ending cylinder. By default, YaST uses the next available starting cylinder for the beginning of the new partition and the last available cylinder as the end (grow to fill). To define the size and location, you can either enter absolute cylinder numbers, or you can use the default start cylinder and enter the size of this partition in kilobytes or megabytes (for example entering `+10M` would create a 10 MB partition).

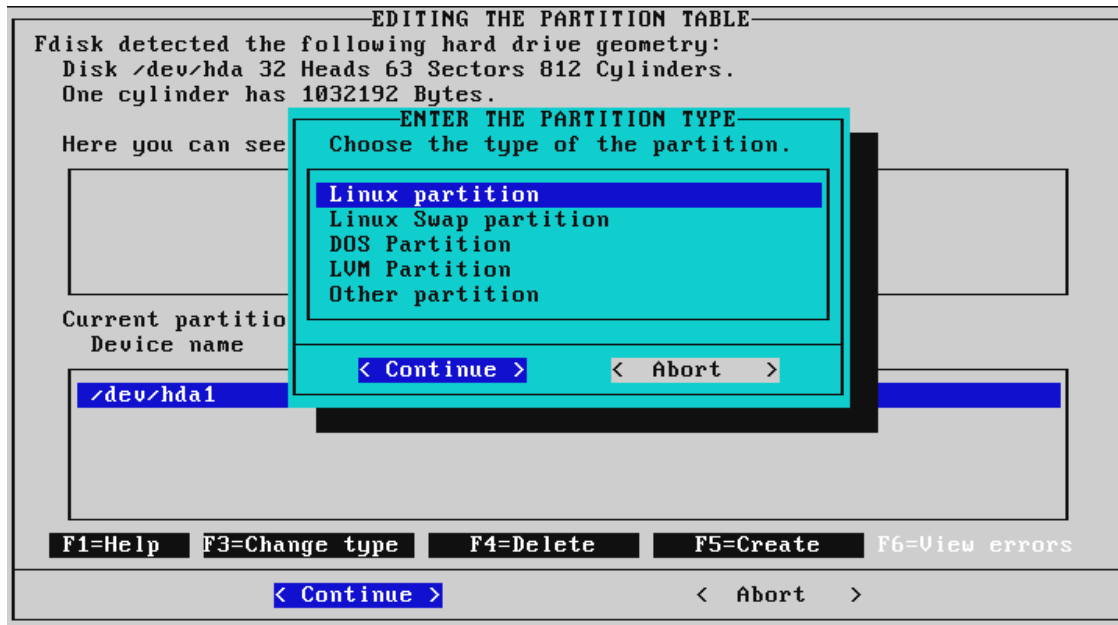


Figure 18. Defining the swap partition

By default, YaST creates Linux native partitions. To create partitions of another type (for example, swap), press F3 after you have selected the desired partition you want to change. Note that this procedure only sets the partition ID of this partition. It does not modify the partition's content or size. Partition the drive(s) to suit your needs.

After the partition table definition is completed, click **Continue** to write the partition table to disk and proceed to the file system creation dialog.

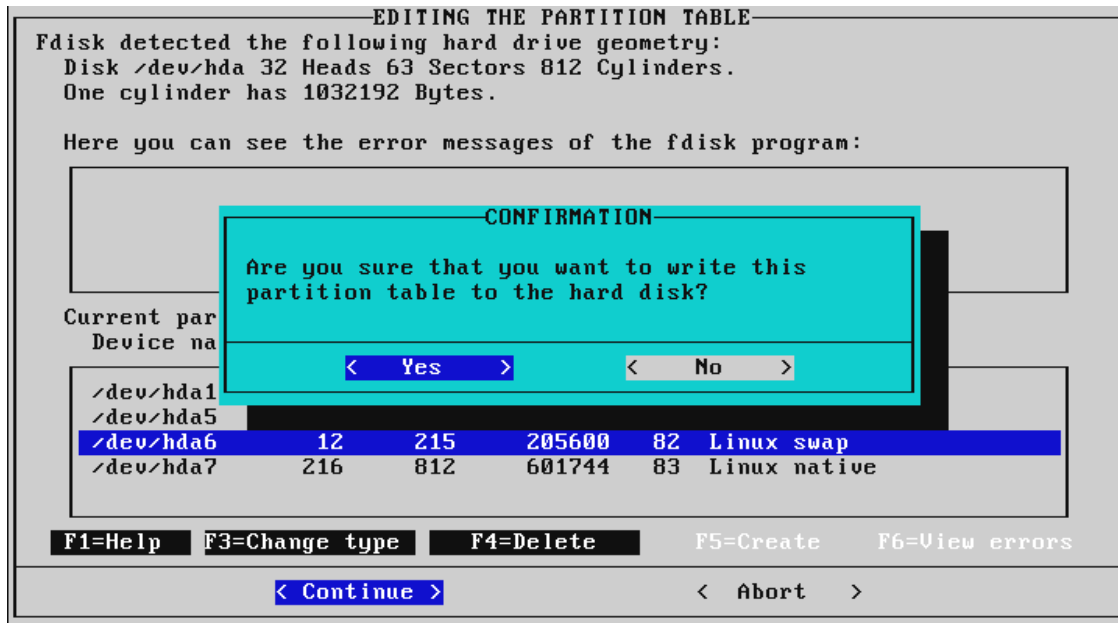


Figure 19. Writing the partition table

Click **Yes** if you want to write the new partition table to this disk. Selecting **No** will abort the partitioning.

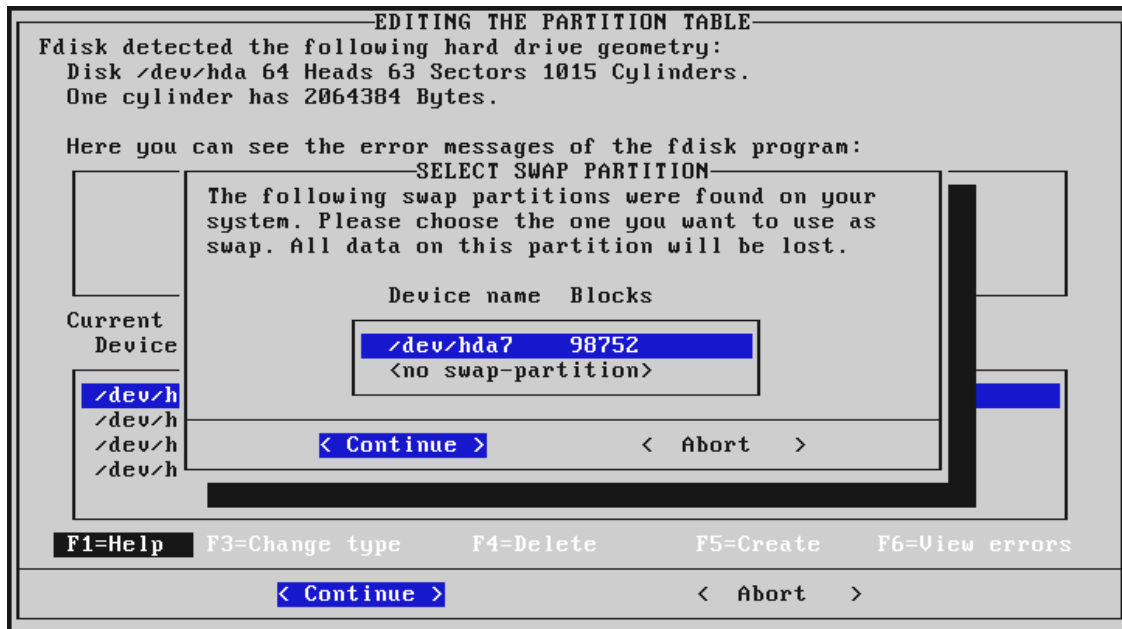


Figure 20. Adding swap space

If you have created a swap partition, YaST will immediately attempt to use it to have more virtual memory for the further installation procedures. Select **Continue** to make use of this. The content of this partition will be deleted, so double check that you are creating the swap space in the correct partition.

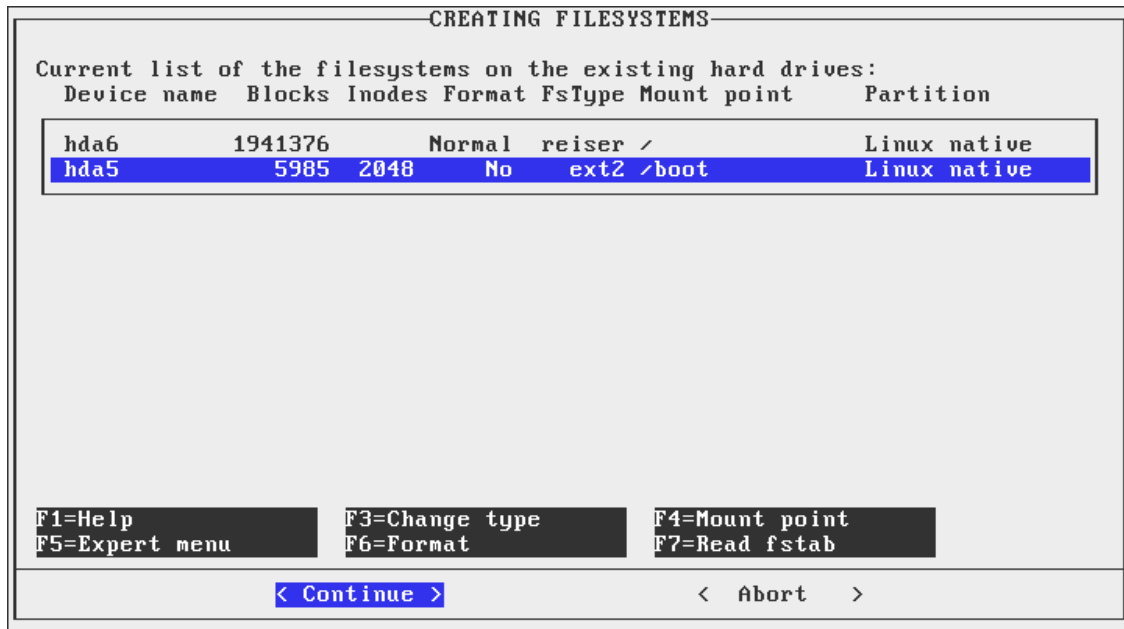


Figure 21. File system creation

After the partition table has been written, you need to create file systems on all partitions that you want to use for Linux (this is similar to formatting them). Additionally you have to define mount points, which is the partition that will act as your root file system, and where other partitions should be mounted. Press F4 to open the mount point dialog.

SuSE Linux provides a new journaling file system that monitors all changes to the system disk (hence the *journaling*) and can play back those changes in the event of a system crash. Users who have experienced a system failure on a 20 GB hard disk running ext2 will know how long it can take for the file system to be checked at boot up for errors. With ReiserFS a 20 GB hard disk can be checked by the system in a matter of seconds. ReiserFS is now considered stable enough to be used in production systems; therefore, you have the option of creating your partitions using ReiserFS. You can see in Figure 21 that we have defined the / (root) mount point as a ReiserFS file system.

Important

Do not set / as a ReiserFS partition unless you have created a /boot partition. The usual way to set up a ReiserFS system is to create a /boot partition formatted as ext2, and then create a / partition formatted as ReiserFS.

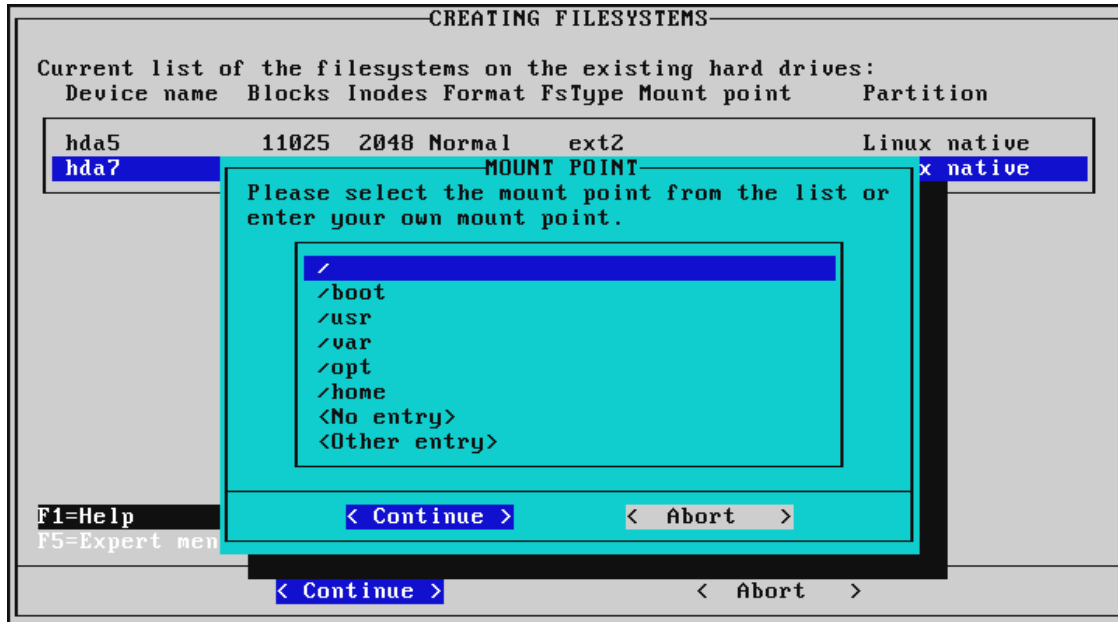


Figure 22. Selection of mount points

YaST offers a list of commonly used mount points. You can either select one from the list or select **Other entry** to define another mount point.

Note

One of your partitions must be mounted to "/". This will be your root partition. YaST will check for the existence of this mount point before proceeding with the installation process.

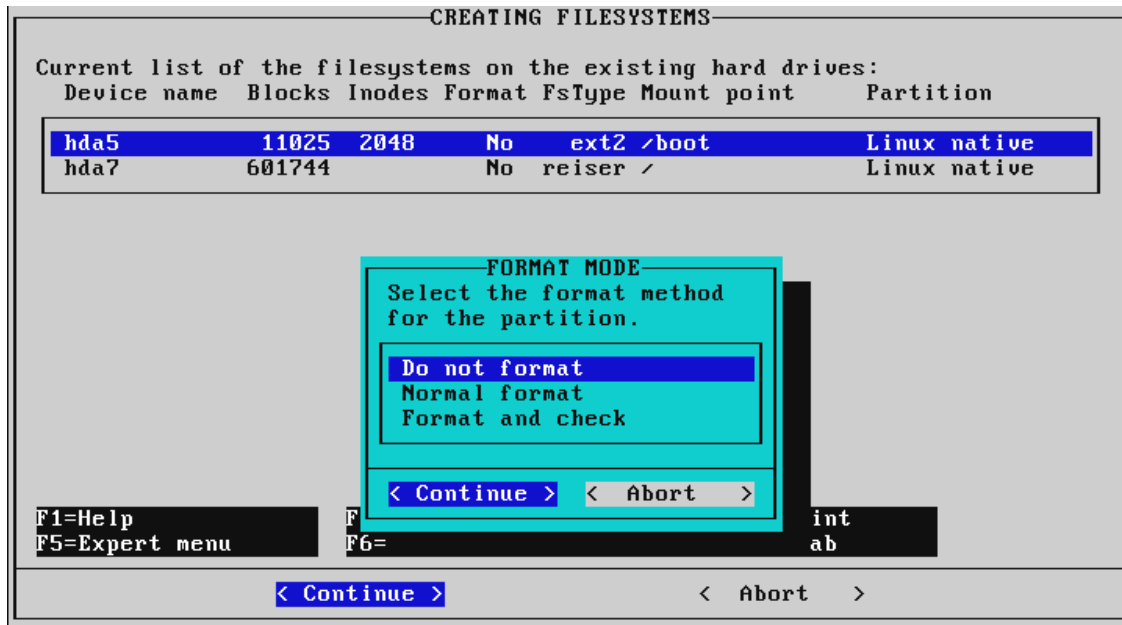


Figure 23. Format mode

After defining the mount points and the type of formatting, select **Continue** to proceed to the creation of these file systems.

Note

This is the same as formatting your hard disk! You will not be able to recover any data that has not been backed up yet! If you are sure, that you want to proceed, select **Yes**.

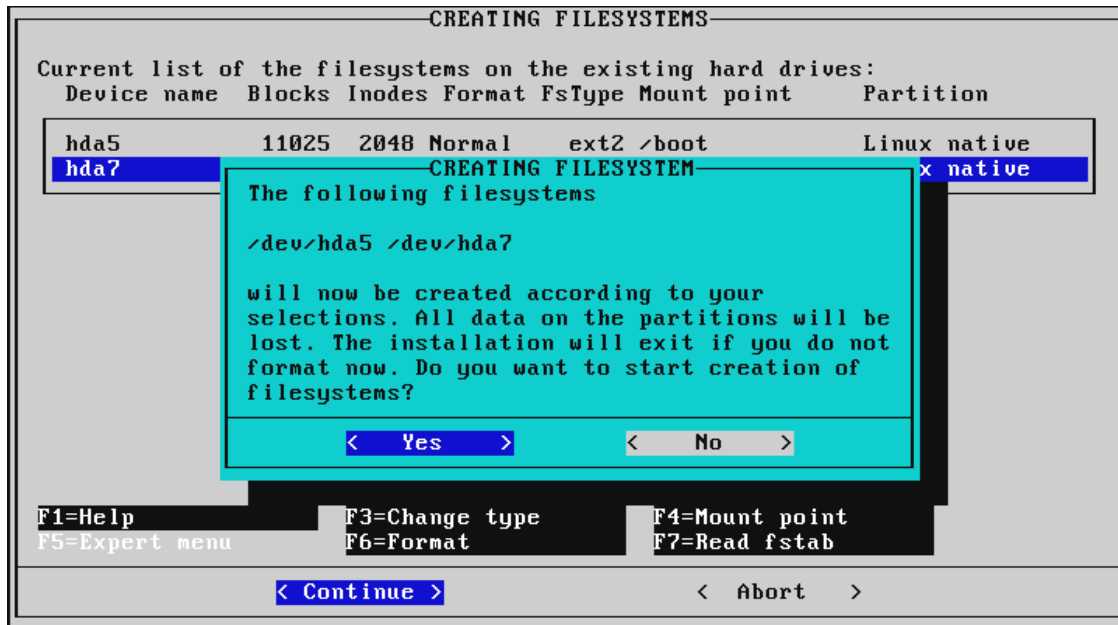


Figure 24. Confirmation to commit your file system settings

The creation of file systems may take some time, depending on the size of your partitions. You should note some hard disk activity during the process.

After the file systems have been successfully created, you will reach YaST's package selection window.

2.3.4 Software package selection and installation

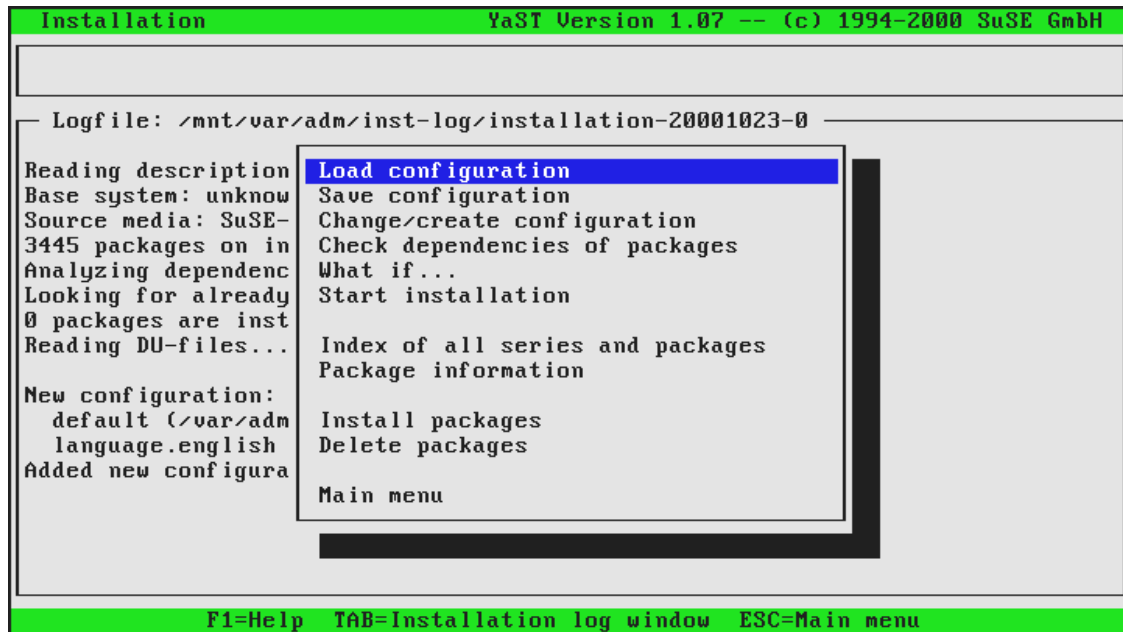


Figure 25. YaST: Package selection

After you have finished the creation of partitions and file systems, the installation system needs to know what software you want to install. Since SuSE Linux offers a broad variety of software packages, it would be a very time-consuming task to check each single package for installation.

The window shown in Figure 25 enables you to define the software packages that will be installed on your system. You will be able to add or remove packages later on; therefore, we will stick with the default configuration here. SuSE also created a number of predefined package selections (configurations), which you can choose with the menu option **Load configuration**.

More information about package management can be found in 3.1 “Adding and removing software packages using YaST” on page 111.

The only package that we will be adding here is the correct driver for the video card, also referred to as the X server.

Select **Change/create configuration**.

```

Series selection YaST Version 1.07 -- (c) 1994-2000 SuSE GmbH
Series
doc Documentation [144.7 M]
e Emacs [ 31.7 M]
emu Emulators [ 0 B]
fun Games and more [ 1.9 M]
gnm GNOME - GNU Network Object Model Environment [ 0 B]
gra All about graphics [ 59.9 M]
ham Amateur Radio (AX.25, CW, Logs, etc.) [ 0 B]
k2de KDE2 - K Desktop Environment (version 2) [ 0 B]
kde K Desktop Environment [ 76.8 M]
kpa KDE applications [ 6.5 M]
n Network-Support (TCP/IP, UUCP, Mail, News) [ 54.2 M]
pay Commercial Software [355.3 M]

<F3>=Zoom
device-name partition total used free free% mount-point
/dev/hda6 Linux 1.85 G 1.12 G 753.3 M 39% /
/dev/hda5 Linux 5.5 M 325.0 K 5.2 M 94% /boot

F1=Help F4=Resorting F5=Dependencies F10=Esc=Exit

```

Figure 26. YaST: Series selection

Figure 26 shows the series selection of YaST. All software packages have been categorized into different series, to make it easier to find the correct program for your needs.

Scroll down and select **xsrv** to open the list of available X servers.




Figure 27. YaST: Package selection

After you have selected a series, you will see a list of all packages available in this series. Select **xsvga**. F2 will give you a more detailed description of the current package. To confirm your selection, press F10 to return to the package series selection menu. You can now select or deselect packages from another series or press F10 to continue.

Now that the selection is finished, you can start the installation of the selected packages by choosing **Start installation**.

```

Installation                               YaST Version 1.07 -- (c) 1994-2000 SuSE GmbH
-----
Installing package 20:      gawk - 691.8 K - 318 packages remaining...

Logfile: /mnt/var/adm/inst-log/installation-20001023-0
at      #####
Postinstall at...
  Updating etc/rc.config...
base    #####
bash    #####
bdflush #####
compress #####
cpio    #####
cracklib #####
cron    #####
ddrescue #####
devs    #####
diff    #####
eazy    #####
ext2fs  #####
file    #####
fileutil #####
find    #####

gawk - GNU awk

```

Figure 28. YaST: Package installation in progress

The installation of software packages from the CD-ROM to your hard disk will begin. Depending on the speed of your CD-ROM and the number of packages, this may take a while. You will be prompted to change the CD-ROM from time to time, to install the remaining packages.




Figure 29. YaST: CD changing prompt

After the installation of packages has completed, YaST will return to the package installation menu shown in Figure 25 on page 29. You are free to add or remove further software packages and to reiterate through this process. To continue the installation of SuSE Linux, select **Main menu**.

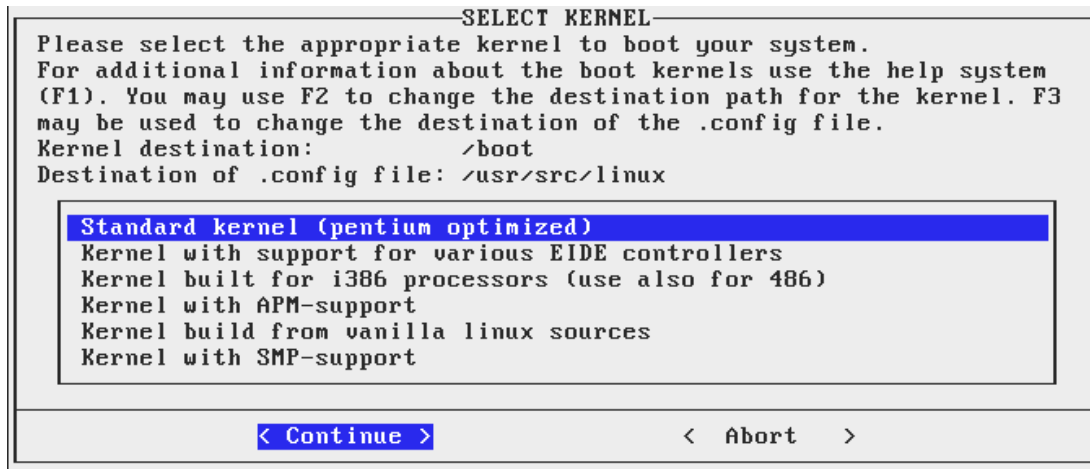


Figure 30. YaST: Kernel selection

In order to be able to boot the installed system from hard disk, you need to install a Linux kernel that is built for your system. SuSE's kernel images are built to be run on many configurations. Choose the kernel image that best fits your needs. If you are using a multiprocessor system, select the **Kernel with SMP-support**. Otherwise, it is safe to select **Standard Kernel**.

You will be prompted as to whether you wish to configure LILO, the Linux Loader. Select **Yes**.

2.3.5 LILO - the Linux Loader

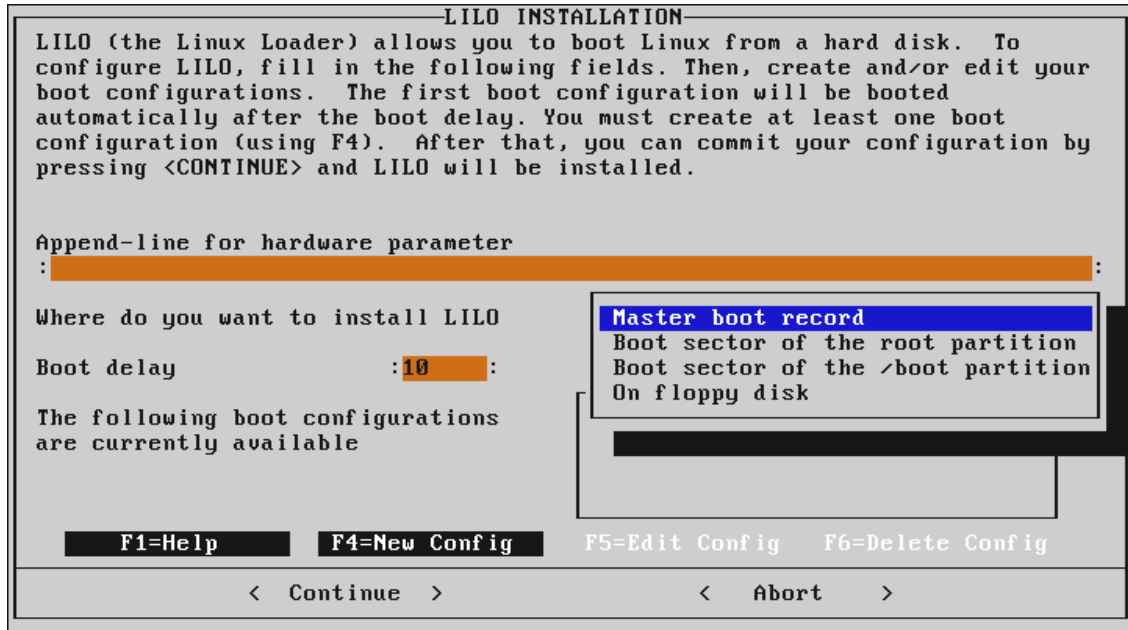


Figure 31. YaST: LILO configuration

LILO, the Linux Loader, is a boot manager that allows you to boot multiple operating systems that can reside on different hard disk partitions or even on different hard disks. Even if you only have Linux installed, you still need to create a boot configuration for Linux. Linux cannot be booted without LILO.

For an exhaustive explanation of LILO and boot concepts, see Chapter 4, “Booting and boot managers: LILO, loadling, etc.” in the SuSE manual.

Figure 31 shows YaST’s LILO main configuration window. You can stick with these default values. However, you might want to decrease the boot delay from 10 seconds (default) to a lower value to save some time during the system startup. By default, LILO will be written to the master boot record of your primary hard disk. Alternatively you can write it to a floppy disk, which has to be inserted during the system bootup.

Press F4 to create a new boot configuration.

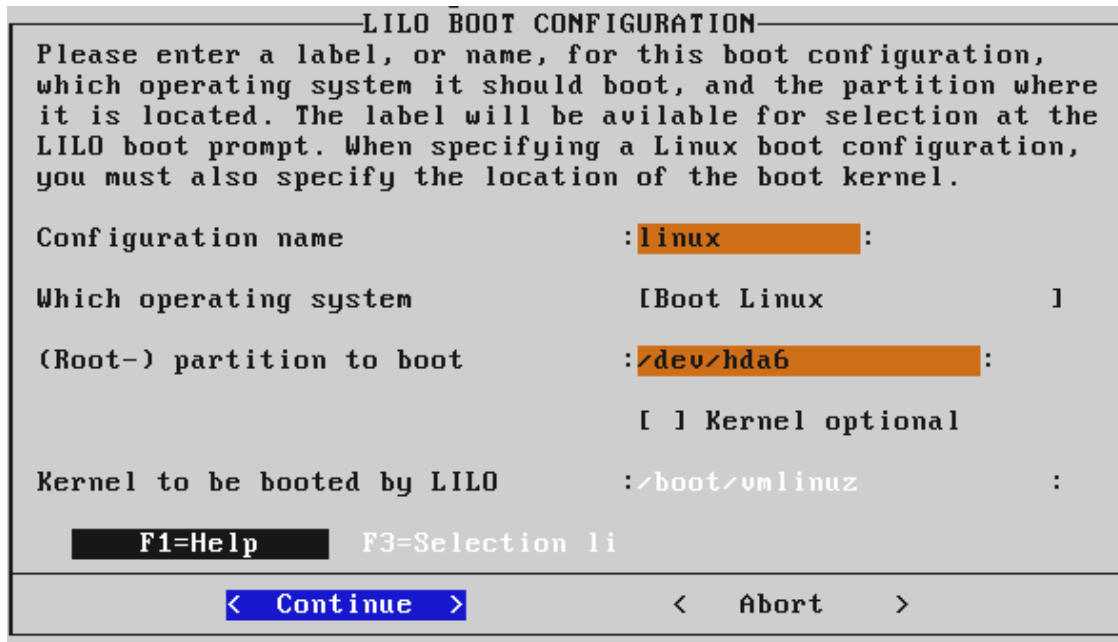


Figure 32. YaST: Create LILO boot configuration

Figure 32 shows the boot configuration dialog. You need to create such a configuration for each operating system you want to boot.

Enter `linux` as the configuration name. This name identifies your boot configuration and it must be unique for each configuration you create. If you want to boot an operating system later on, you have to enter this name at the LILO: prompt.

In the **Which operating system** field, enter `Boot Linux`. Choose the relevant value if you wish to boot another operating system.

In the (Root-) partition to boot field, enter `/dev/hda6`, where 6 is the partition number. You have to select your root-partition here (the partition that is mounted to `/`) - not the Boot-Partition! Usually this is already correctly preselected.

In the Kernel to be booted by LILO field, enter `/boot/vmlinuz`. LILO needs to know where the kernel image to be booted is located.

When you have added all necessary boot configurations, select **Continue** in the LILO main menu to write the new boot record.

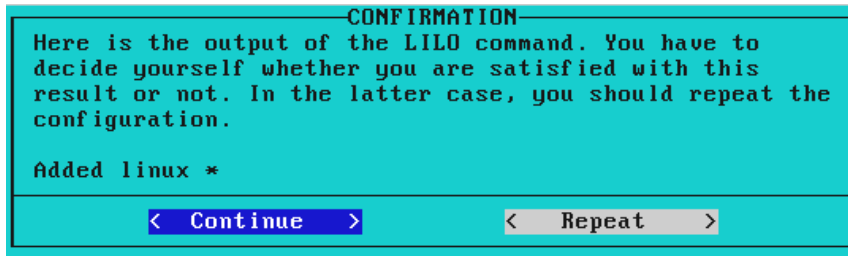


Figure 33. YaST: LILO output

After LILO has created the new boot block, you are given the opportunity to review LILO's output. Click **Continue** to proceed to the time zone and clock settings.

2.3.6 Time zone and clock settings

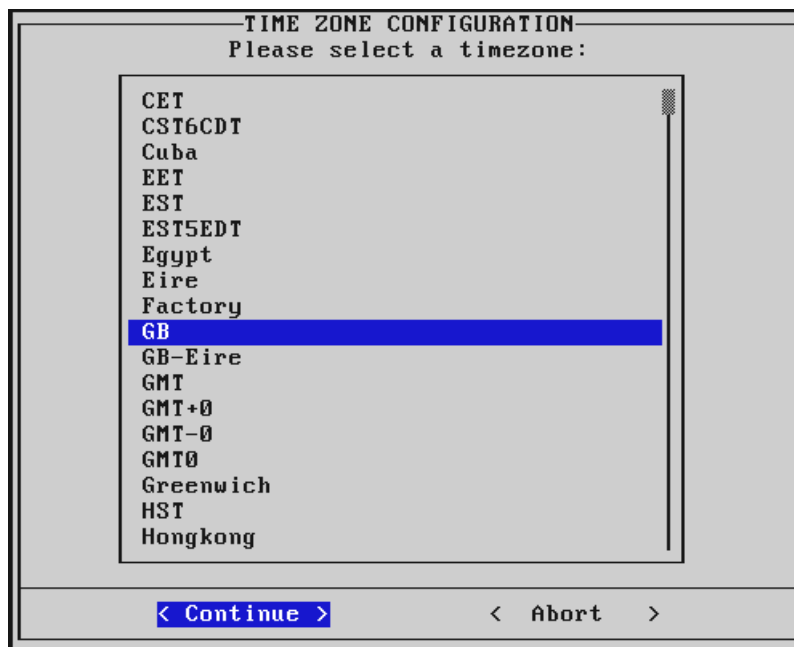


Figure 34. YaST: time zone configuration

Select your desired time zone here. This is important for automatic switching between summer and winter time. You can also change this value later on, which is helpful if you use your Linux system in different locations (for example on a laptop computer).




Figure 35. YaST: system clock selection

The setting of your BIOS clock has to be selected. Click **Local time** if you are booting other operating systems on this box; otherwise choose **GMT**.

2.3.7 Network configuration

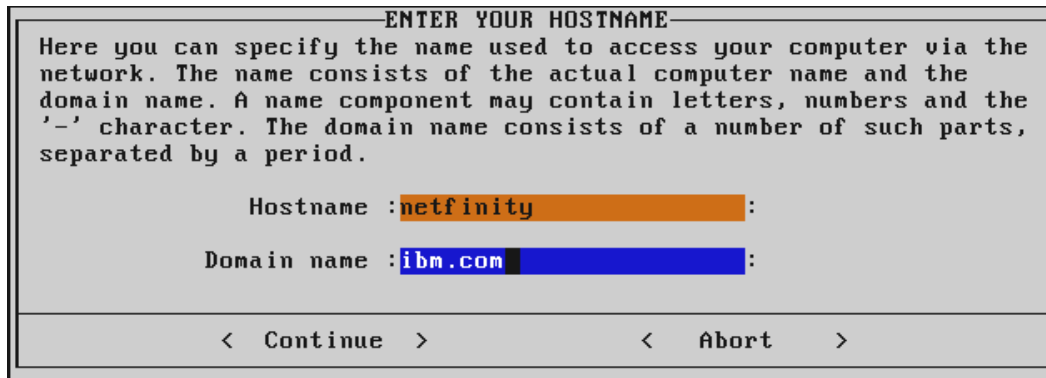


Figure 36. YaST: host name configuration

Enter your host and domain name here. Each host in a TCP/IP network must have a unique host name. If you do not know this, contact your network administrator of your local network. If you do not intend to use this system in a networked environment, you can choose your host and domain name.

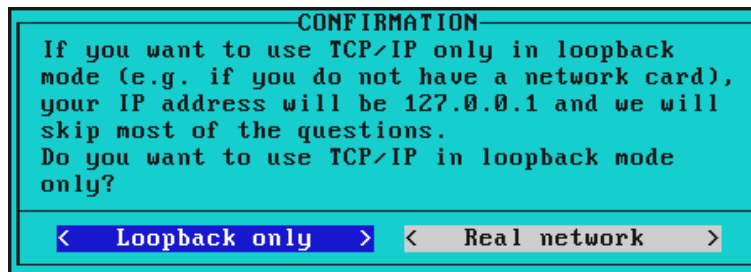


Figure 37. YaST: network type selection

If your system will be connected to a LAN, select **Real network**.

Otherwise, select **Loopback only**. This will skip the following questions and the installation will continue with the Sendmail configuration shown in Figure 40 on page 40.

For DHCP client selection, select **No**, if you will use a static IP address for the network card (which is recommended for a server). If you select **Yes** here, the system will act as a DHCP client in your network and will obtain its IP address from a DHCP server in your local network. In this case, the window shown in Figure 38 will not appear.

```
-----ENTER THE NETWORK ADDRESSES-----
Please enter the data required for the configuration of your
network. These are the IP address you want to give the machine
currently being installed (e.g. 192.168.17.42) and the netmask of
your network. The latter is 255.255.255.0 for most of the (smaller)
networks, but you may wish to set it to a different value. If you
need a gateway to access the NFS server, please enter the IP
address of the gateway host.

Type of network: [eth0] ]
IP address of your machine: [192.168.0.99] :
Netmask (usually 255.255.255.0): [255.255.255.0] :
Default gateway address (if required): [192.168.0.8] :
IP address of the Point-to-Point partner :

[ < Continue > ] [ < Abort > ]
```

Figure 38. YaST: network configuration

The window shown in Figure 38 enables you to configure your TCP/IP configuration. You need an IP address to be able to communicate with other hosts in your network. Contact your network administrator for the correct values for your network.

Type of network - select the desired network card here. Select **eth0** to use the first Ethernet card, **tr0** if you use a token-ring adapter.

Enter the correct values for your local network and click **Continue**.

You will now be prompted if you want to start the inetd service. Inetd is needed for invoking certain services on demand, such as telnet, finger, ftp

and others. Inetd should always be started; otherwise, the above-mentioned services will not be available. If your system will be connected to the Internet, you may want to restrict access to certain services. Please see section 18.2.2, “inetd” in the SuSE manual for more information about inetd. In most cases it is safe to select **Yes** here.

If you want to use this system as an NFS or NIS server, you will need to start the portmapper service at boot-up. Therefore, the question “START THE PORTMAPPER?” should be answered with **Yes**.

If you have decided to start the portmapper, you will now be prompted as to whether you want to start the NFS server as well. Select **Yes** if you plan to share files using NFS.

The ADJUST NEWS FROM-ADDRESS dialog window enables you to modify the sender address, if you intend to use Usenet News. The default is fine here for most cases; select **Continue** to proceed.

If your system is connected to a network and you would like to access a Domain Name System (DNS) server, select **Yes** at the confirmation (nameserver) dialog. If your system will act as the DNS, select **Yes**, too.

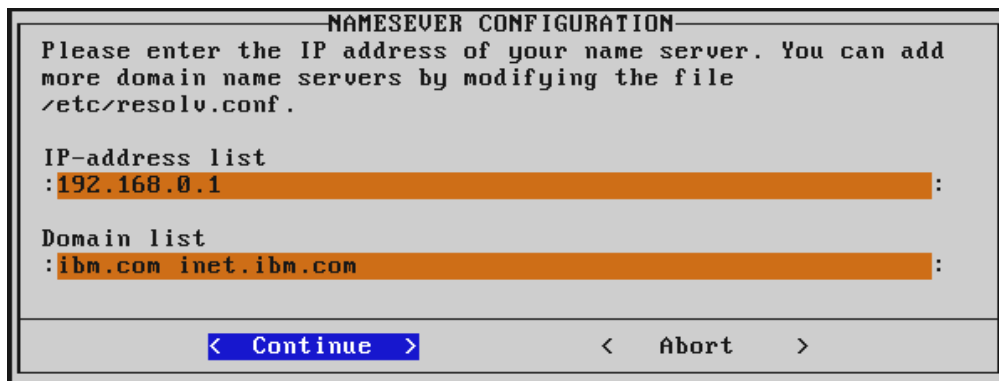


Figure 39. YaST: nameserver configuration

Figure 39 shows the name server configuration dialog. You can enter your name server’s IP address on the first line. If you want to access multiple name servers, separate the entries with a space. Adjust the domain list to your local domain.

If you want to run a DNS server on this system, you still have to configure the system to query the local running name server. Select the loopback interface (127.0.0.1) as the name server’s IP address.

Choose **Continue** to advance to the next window.

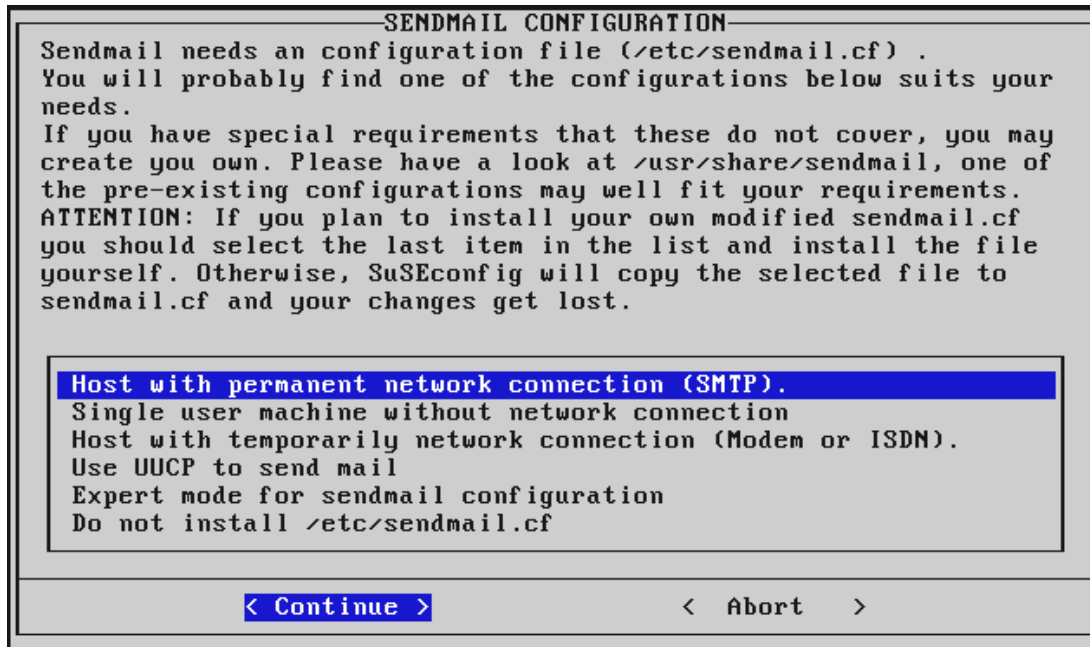


Figure 40. YaST: Sendmail configuration selection

You will now be asked how you want to install the Sendmail service. The default selection is good for most configurations. Press **Continue** after you have made your choice.

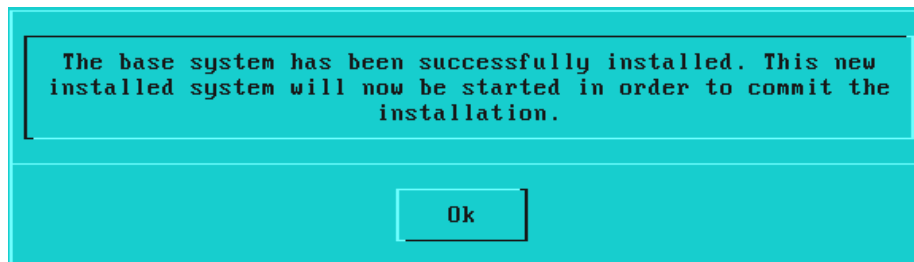


Figure 41. Completion of package installation

After you have entered all the necessary values, the installed system will be booted up.




Figure 42. Definition of the root password

The installation program will now start to boot up from the freshly installed system. Since Linux is a multi-user operating system, you have to define user accounts first. The most important user account is the root account, which identifies the system administrator (username “root”) of this system. Each user account is protected by a password. Therefore, you will be prompted to enter a password for the root user twice. Please note that passwords in Linux are case-sensitive!

The next window will ask you to set up your modem. If you have one, click **Yes**. Clicking **No** will skip the following window.




Figure 43. YaST: modem configuration

Figure 43 displays the modem configuration window. Select the serial port to which your modem is connected. YaST will create a symbolic link /dev/modem that will point to the respective serial device. Please note that this is only the first step in configuring your modem for Linux. The symbolic link just makes it easier for other applications to find the modem. However, these applications

still have to be configured manually to be able to “talk” with the modem later on. Click **Continue** to create the link.

2.3.8 Mouse configuration

After configuring your modem, you can now configure the mouse. If you intend to use the X-Windows system later on or want to use the mouse on the text console, click **Yes**. If you do not need a mouse, click **No** to skip the following mouse configuration dialogs.

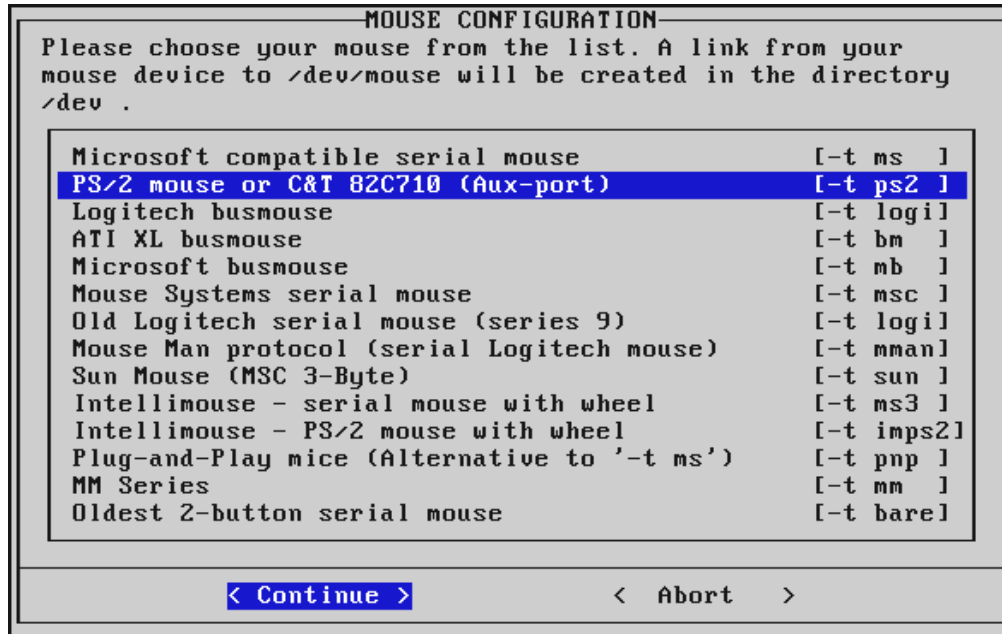


Figure 44. YaST: mouse configuration

First, you have to choose the type of mouse you have. The two most common types are Microsoft-compatible or PS/2 mouse. IBM Netfinity servers use PS/2; therefore, select **PS/2 mouse**. If your mouse is connected to a serial port, it is most likely a Microsoft-compatible mouse. If you choose a serial mouse, you will also have to select the correct serial port as shown on Figure 45.

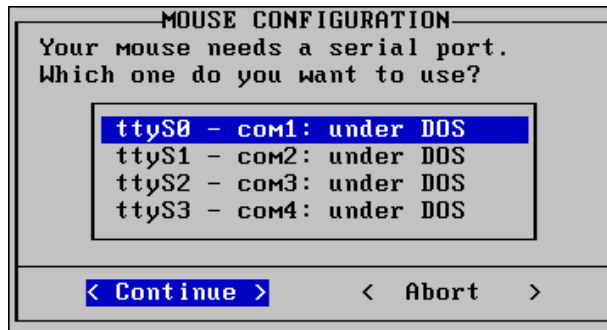


Figure 45. YaST: serial mouse port selection

YaST will create a symbolic link `/dev/mouse`, which will point to the correct mouse device (for example `/dev/psaux` for PS/2 mice or `/dev/ttyS0` for a serial mouse on the first serial port).

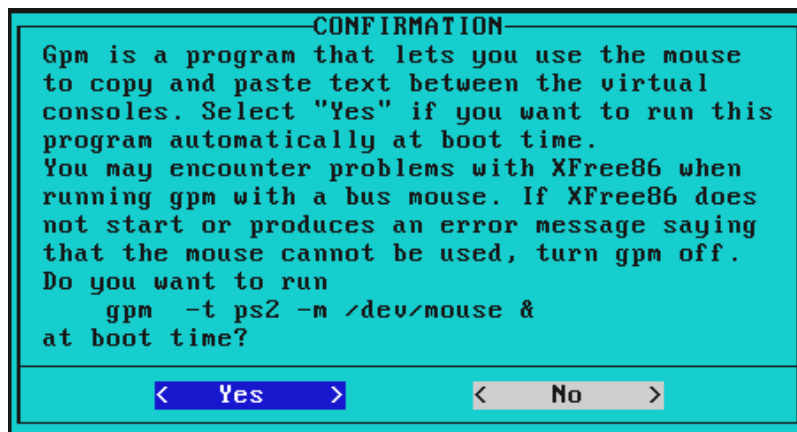


Figure 46. YaST: GPM configuration

General Purpose Mouse (GPM) is a helpful program, if you do a lot of work on the command line in text mode. It enables you to copy and paste text between virtual consoles by highlighting the text with the mouse. Some applications, such as the Midnight Commander (MC), can also be operated with the mouse. Select **Yes** if you want GPM to be started on system startup. Selecting **No** will skip the following window.

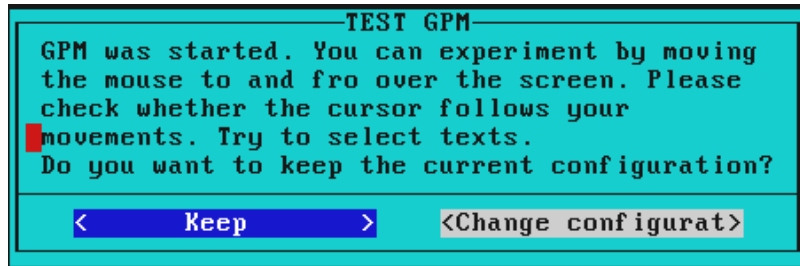


Figure 47. YaST: GPM test window

YaST will now start GPM to let you test your configuration. Try to move the mouse around; the cursor should follow your mouse movement. Also try to select some text by highlighting it with the left mouse button. If the cursor does not move at all, or jumps randomly across the screen, you have most likely chosen the wrong mouse protocol. Click **Change configuration** to return to the previous window and try another mouse protocol. If everything is working correctly, click **Keep** to continue. Unfortunately you cannot use the mouse for this.

YaST will now terminate to boot the system.

```
Setting up network device eth0 done
Setting up routing (using /etc/route.conf) done
Re-Starting RPC portmap daemon done
Re-Starting syslog services done
Loading keymap qwerty/us.map.gz done
Initializing random number generator done
Starting kernel based NFS server done
Starting service httpd done
Starting service at daemon: done
Starting console mouse support (gpm): done
Starting INET services (inetd) done
Starting lpd done
Initializing SMTP port. (sendmail) done
Starting CRON daemon done
Starting Name Service Cache Daemon done
Master Resource Control: runlevel 2 has been reached
Please enter "root" to login as user root...

Welcome to SuSE Linux 7.0 (i386) - Kernel 2.2.16 (tty1).

netfinity login:
```

Figure 48. SuSE Linux login

Log in as user root with the password you provided during the installation to finalize the installation. You can also log in and start working with the regular user account you have created.

The basic installation of SuSE Linux is now completed. You can start configuring the X-Windows system (2.6 “XFree86 configuration” on page 103) and the additional services.

2.4 Installation using Yast2

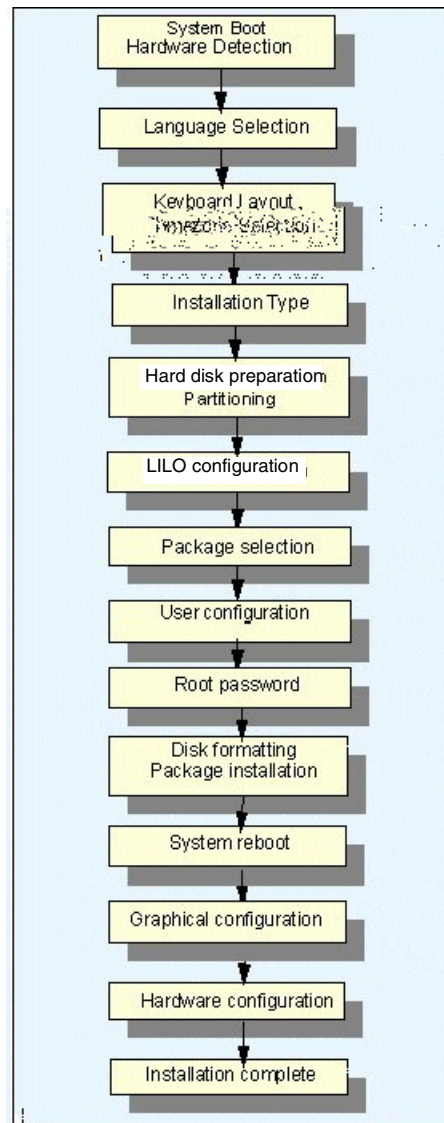


Figure 49. Yast2 installation process

To start installing SuSE Linux with Yast2, insert the first CD into the CD-ROM drive and power on the machine. After a few seconds you should see the window in Figure 50. If you do not, make sure you have configured your boot parameters correctly to boot from either the CD-ROM or floppy drive.

2.4.1 Yast2: Booting the installation system



Figure 50. Booting the SuSE Installation

The installation will load the kernel and autodetect the system hardware. After the kernel and installer have been loaded, you will be presented with the language selection window in Figure 51.

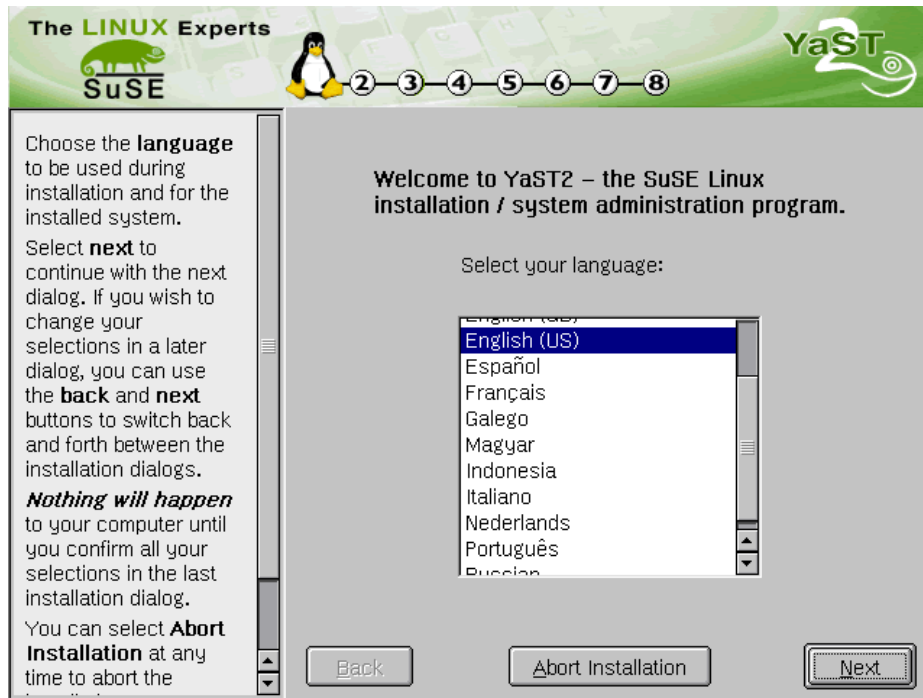


Figure 51. Language selection

Note

In the left hand pane of Yast2 there are instructions on how to proceed through the current window. If these instructions do not fulfill your needs, look in the user manual for a more thorough explanation.

Select your desired language from the list and click **Next** to continue. The language you select on this window dictates the language the installation will continue in and also the language of the installed system.

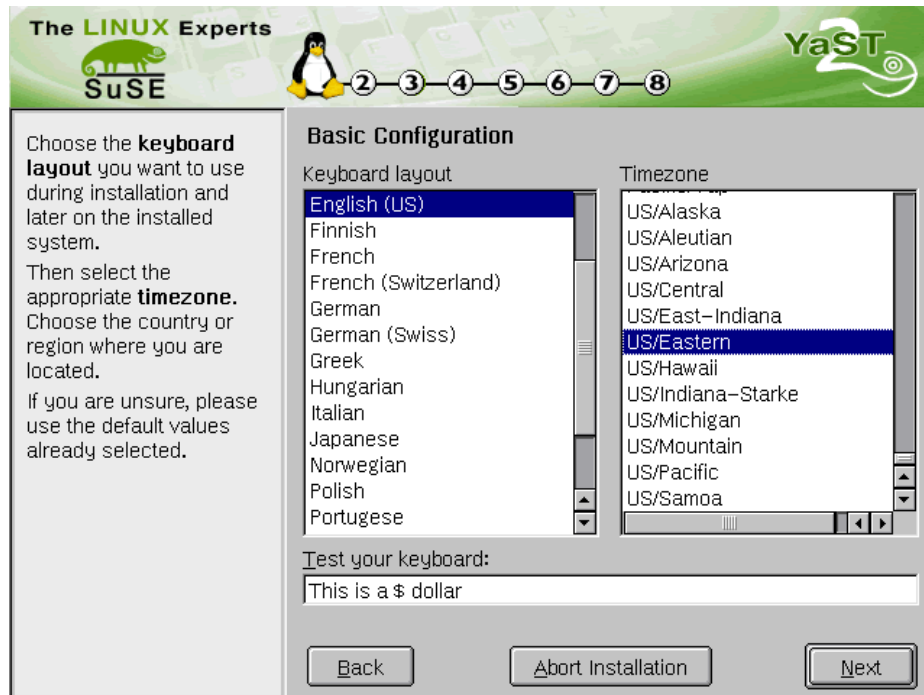


Figure 52. Keyboard and time zone settings

Select the desired keyboard layout and time zone. It is important that you choose the correct keyboard settings from this window, since it can be very difficult to enter data into the system if you choose an incorrect keyboard map. If you are in doubt about your keyboard selection, try testing the keymap out in the test area.

2.4.2 Yast2: Installation type

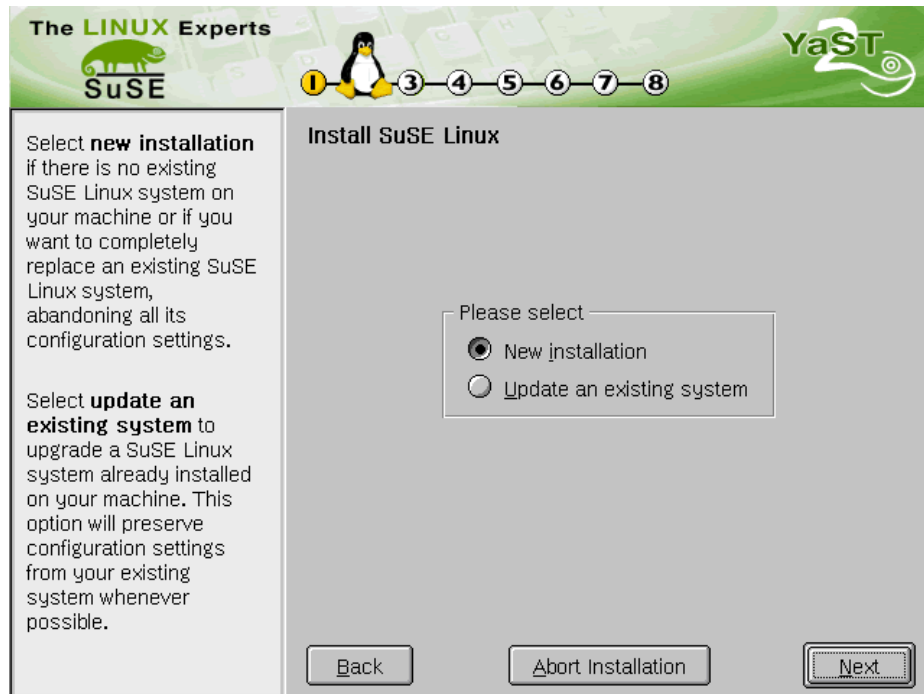


Figure 53. Installation type

Depending on whether you have a previous installation of SuSE, you can choose to either upgrade an existing system or proceed with a fresh installation. SuSE can upgrade applications and services on an existing system to allow a simple upgrade path.

2.4.3 Yast2: Preparing the hard disk

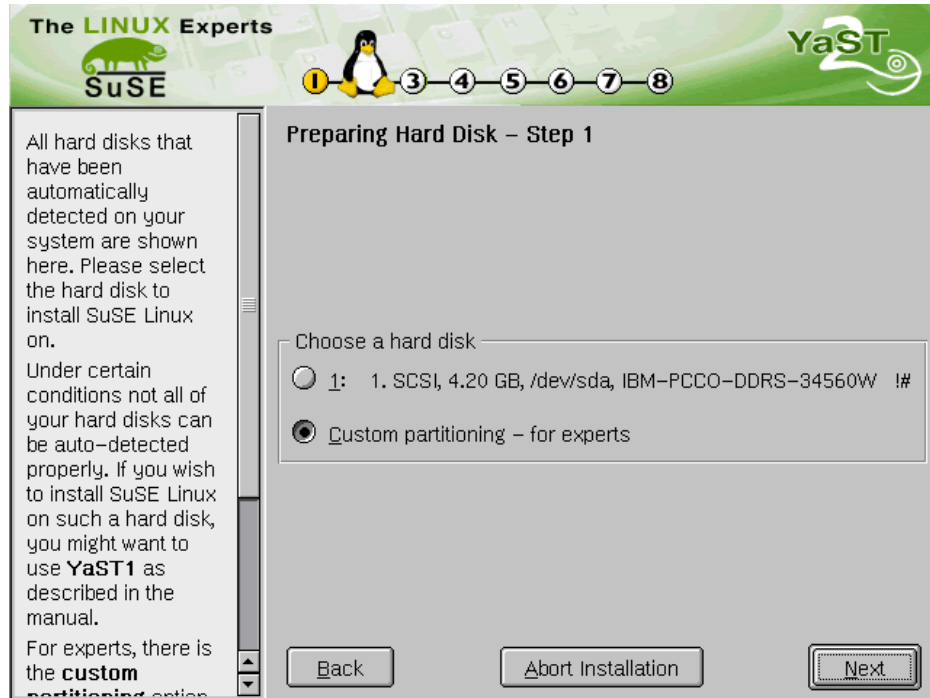


Figure 54. Partitioning method

The SuSE installer can automatically partition your hard drive based on its assumptions of your system. The other option is to partition your system manually, allowing more control over the process.

We will continue with a manual partition setup, since not every system is designed for the same purpose and the default partitioning may not provide an ample file system layout.

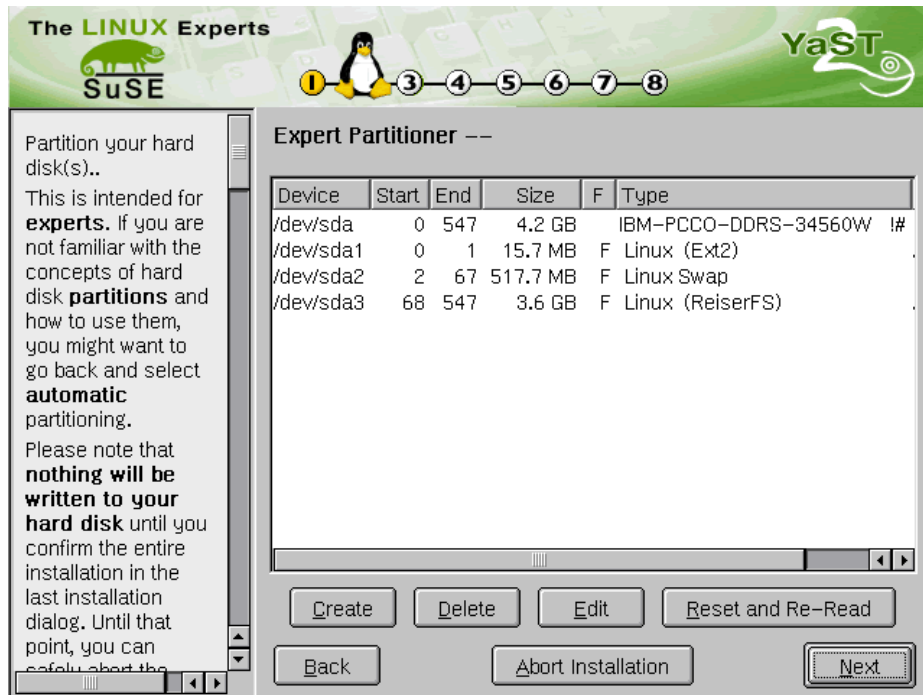


Figure 55. Partition layout

The Expert Partitioner window allows you to create and delete partitions on the hard disks in your system. You have the following options:

- **Create** -This allows you to create a partition on the selected hard drive. After clicking this button you are presented with Figure 56.



Figure 56. Partition type

Selecting **Primary partition** creates a partition on your hard disk that is visible to the partition table. This is the normal partition type selected. Selecting **Extended partition** creates a partition that is also seen by the

partition table, but it allows you to use an extended partition, which helps to work around the four-partition limitation on hard disks.

- **Delete** - This allows you to delete the selected partition.
- **Edit** - Editing a partition entry allows you to modify the format type (Ext2 or Reiser), the size of the partition and the mount point if the partition once it has been created. You can see in Figure 55 what the edit window looks like.
- **Reset and re-read** - This will reset your changes and revert back to the partition table layout that was originally read when the expert partitioner module was loaded.

It is advisable that you create a swap partition for your system. A swap partition allows Linux to “swap out” unused memory segments from main memory to the hard drive, creating more available, faster main memory. The recommended size of a swap partition is between the amount of system memory you have and 1.5 times the amount of system memory you have. If you envisage your system using more swap space, then increase this value. We recommend buying more RAM to suit your memory needs rather than relying on swap memory to provide memory to the system, as swap space is nowhere near as fast as main memory.

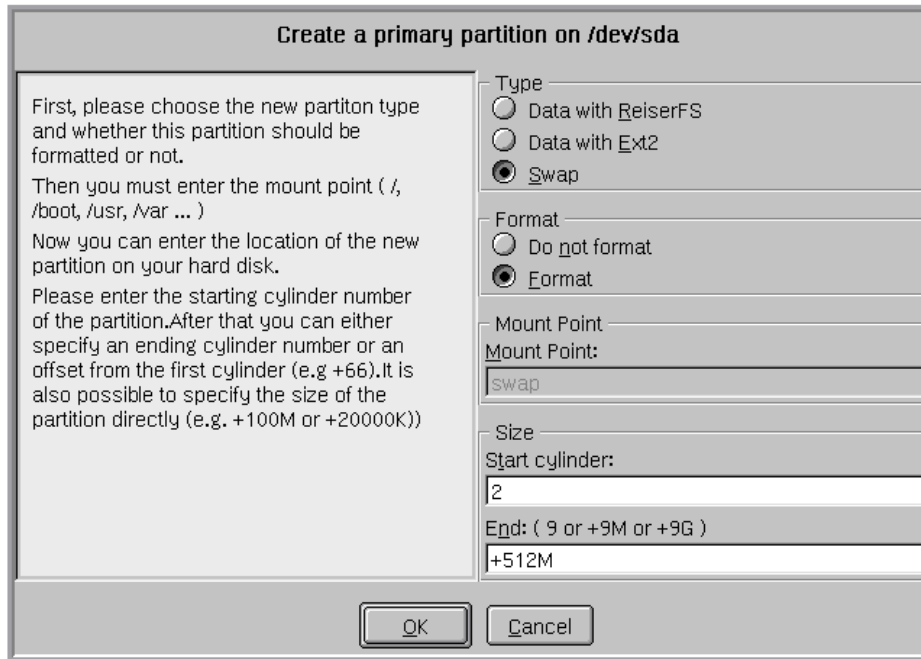


Figure 57. Allocating swap space

SuSE Linux provides a new journaling file system that monitors all changes to the system disk (hence the *journaling*) and can play back those changes in the event of a system crash. Users who have experienced a system failure on a 20 GB hard disk running ext2 will know how long it can take for the file system to be checked at bootup for errors. With ReiserFS a 20 GB hard disk can be checked by the system in a matter of seconds. ReiserFS is now considered stable enough to be used in production systems; therefore, you have the option of creating your partitions using ReiserFS.

Important

Do not set / as a ReiserFS partition unless you have created a /boot partition. The files in the /boot directory have to be read at system boot up and as such should be on an ext2 file system. The usual way to set up a ReiserFS system is to create a /boot partition formatted as ext2, and then create a / partition formatted as ReiserFS.

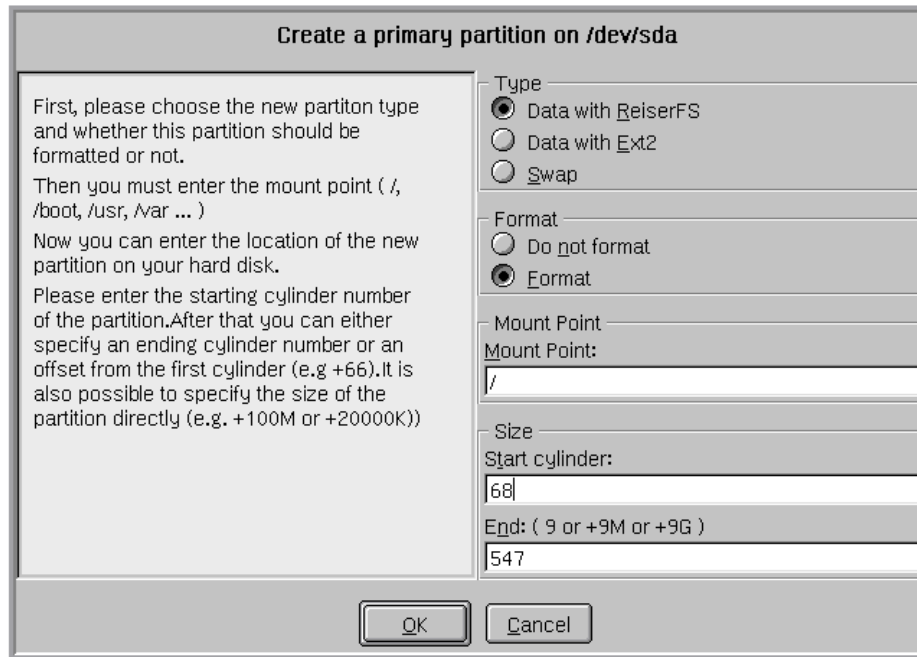


Figure 58. Choosing ReiserFS as your root file system

Once you have created your partitions and file systems, click the **Next** button (see Figure 55). You will be presented with the window in Figure 59.

2.4.4 Yast2: Software installation type

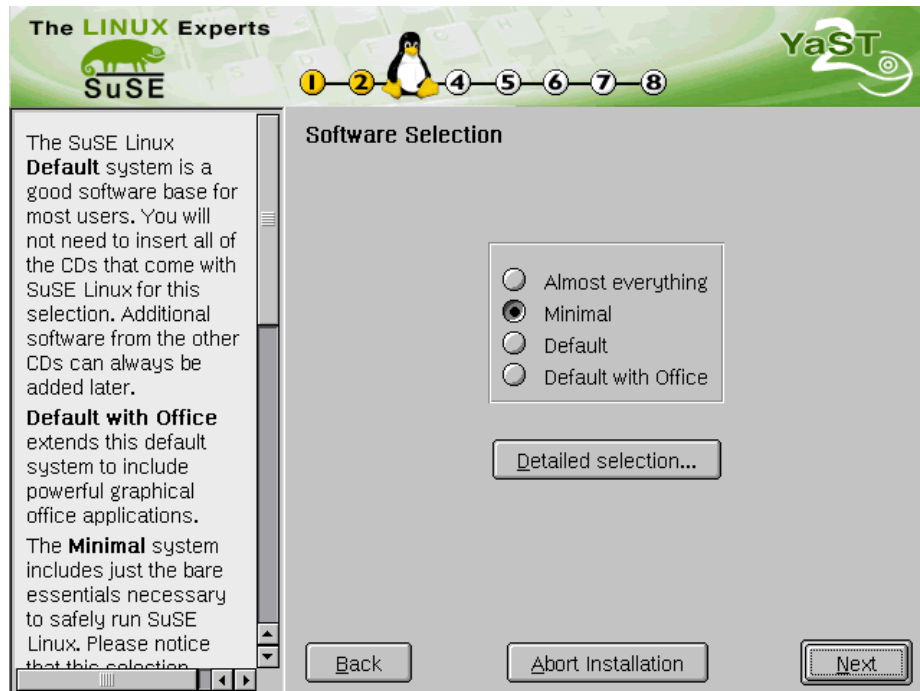


Figure 59. System installation type

Depending on the application of the system you configuring, you can select one of the following options for the installation type:

- **Almost Everything** - This is usually selected by users who wish to try every application in the SuSE distribution. Only the bravest of people should attempt this.
- **Minimal** - This installs only a bare SuSE system. This is ideal for servers, since you can then build upon the base system installed. It guarantees that the services and applications you want for your system are installed at your convenience.
- **Default** - This option installs the default SuSE system. This includes X-Windows and applications used in a working system.
- **Default with Office** - This installs the default system plus applications used for office work. This includes Star Office, the X Office system.

Clicking the **Detailed Selection** button allows you to individually select and deselect applications based on the installation type. If you wish to only install

applications and services that your system needs, then this is the button to click on.

Once you have selected the applications you need in the system, click the **Next** button to continue with the installation.

2.4.5 Yast2: LILO configuration

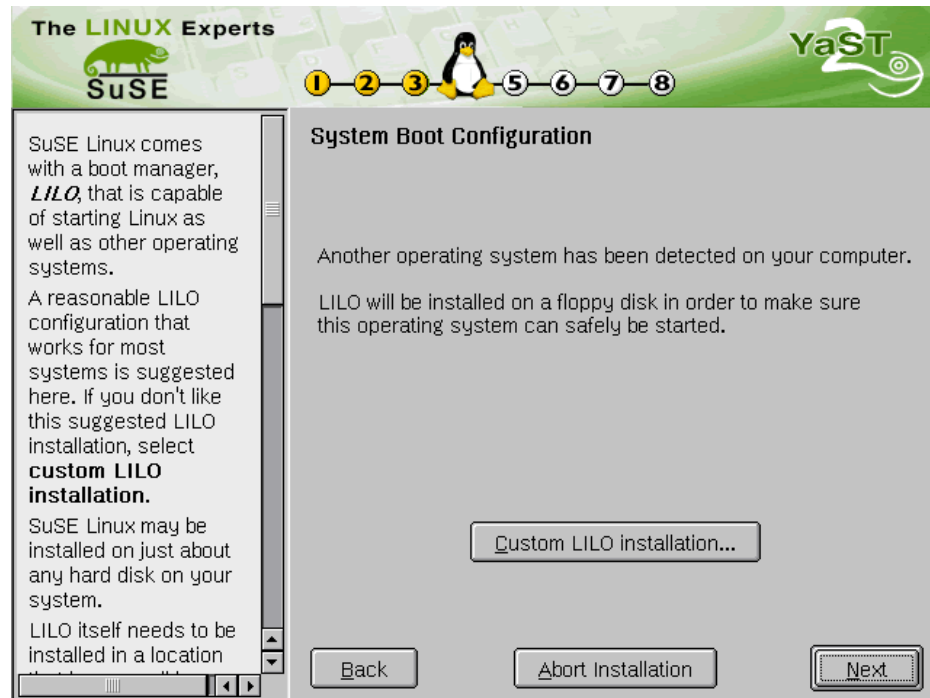


Figure 60. LILO setup

LILO allows your system to load the Kernel image into memory and execute it when the system boots. If you have another operating system on your machine, you will be able to configure LILO to allow you the choice of booting to either Linux or the other operating system you have installed.

If the installation system has not detected an operating system during initialization, you will see Figure 60.

Clicking **Customize LILO Configuration**, you will see Figure 61.

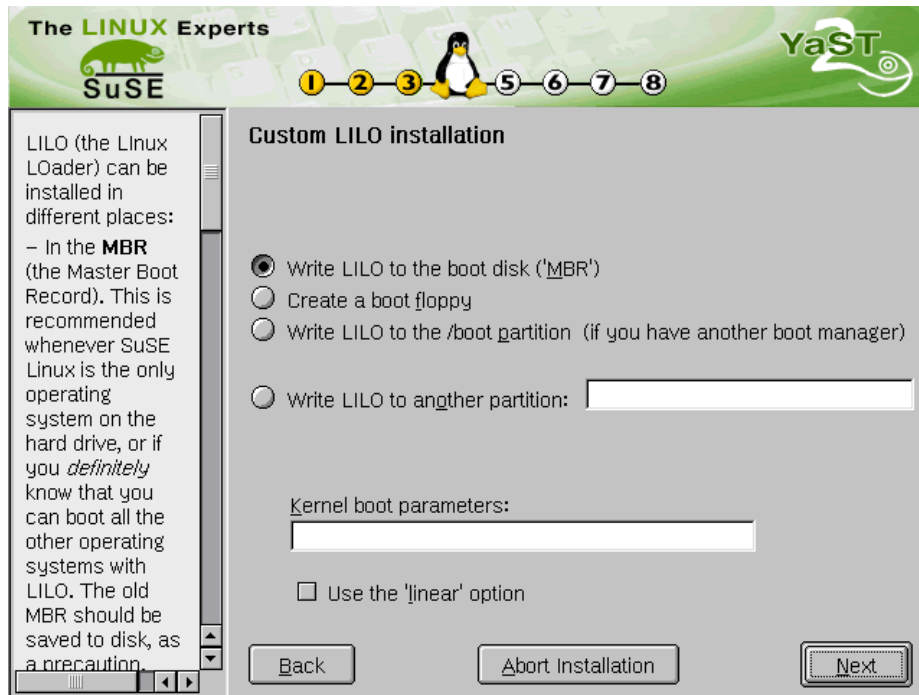


Figure 61. LILO advanced configuration

You can configure where you wish LILO to be installed:

- **Written to MBR** - This creates the LILO configuration in the Master Boot Record of your hard drive. This is the most common option.
- **Create Boot Floppy** - This will allow you to create a floppy disk that can be used to boot your system. Your Master Boot Record will not be modified. This is useful if you wish to avoid your virus scanner complaining about the boot sector being changed, or you wish to add extra security to your system.
- **Write LILO to /boot** - If you have another boot manager in use and do not wish to overwrite it, then this option allows you to configure your current boot manager to load Linux as well as the other operating system.

Note

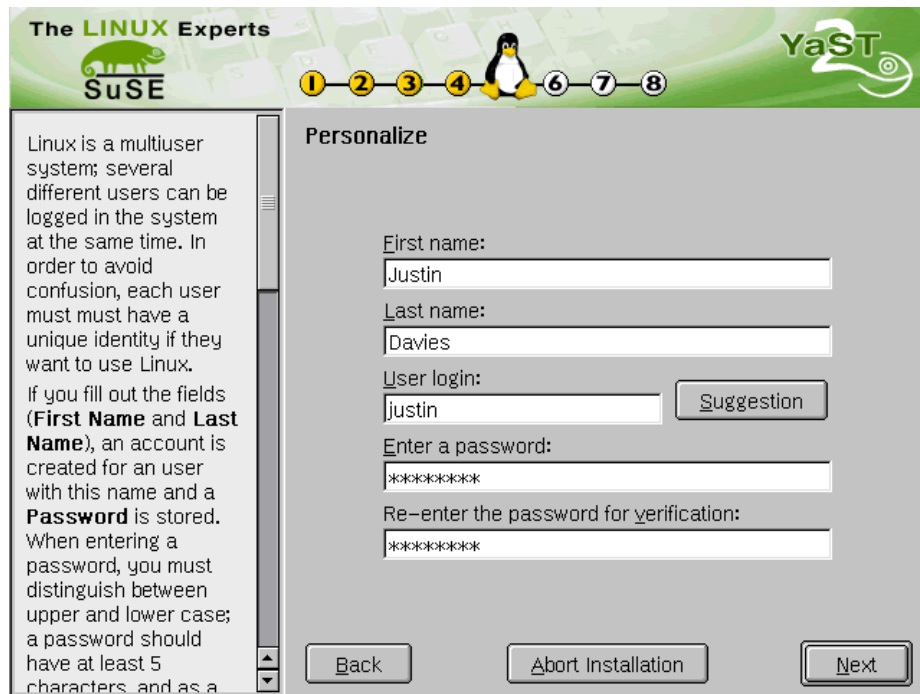
For information on how to configure your boot manager, read the mini-howtos on and Linux co-existing with other operating systems:

<http://www.linux.org/docs/ldp/howto/HOWTO-INDEX/mini.html>

- **Kernel Boot Parameters** - This option allows you to pass the kernel parameters for drivers built into the system. This can be useful for changing the behavior of IDE CD-Writers, or IO settings for network cards.

Click **Next** to continue the installation.

2.4.6 Yast2: Adding a normal user



The screenshot shows the YaST2 installation interface. At the top, it says "The LINUX Experts" and "SuSE" with a logo. A progress bar shows steps 1 through 8, with a penguin icon at step 4. The main window is titled "Personalize" and contains the following fields:

- First name: Justin
- Last name: Davies
- User login: Justin (with a "Suggestion" button next to it)
- Enter a password: [masked with asterisks]
- Re-enter the password for verification: [masked with asterisks]

At the bottom, there are three buttons: "Back", "Abort Installation", and "Next". On the left side of the window, there is a scrollable text area with the following text:

Linux is a multiuser system; several different users can be logged in the system at the same time. In order to avoid confusion, each user must have a unique identity if they want to use Linux. If you fill out the fields (**First Name** and **Last Name**), an account is created for an user with this name and a **Password** is stored. When entering a password, you must distinguish between upper and lower case; a password should have at least 5 characters, and as a

Figure 62. Adding a new user

You will be asked to create a “normal” user for the system. This is a security procedure, since some people have the habit of using a system as the *root* user. This is not advisable since it is very easy to erase a file or directory from the system as *root* if you are not careful. Becoming *root* should only be used when system administration needs to be conducted.

Click **Next** to create the “normal” user and to proceed to setting the root password.

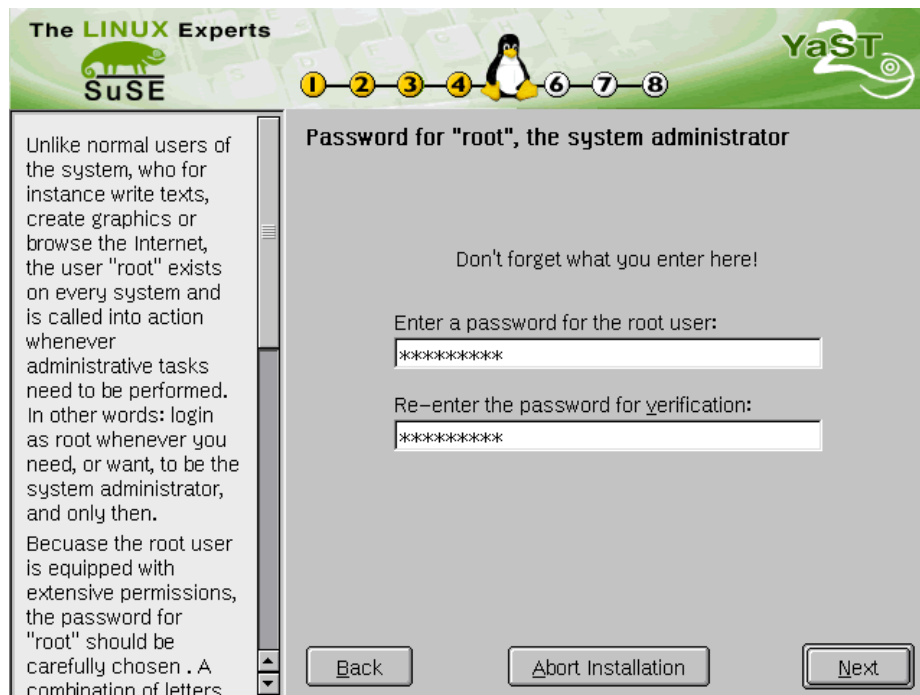


Figure 63. Setting the root password

The root password is the most critical password in the system. It means the difference between a safe controlled system and a compromised system. Therefore, choose a password that is not easy to guess. This means do not use the name of a spouse, a birth date, or a social security number, or anything else that can be guessed from your life.

Use a combination of letters and numbers to create the password, but make sure you can remember it. It is very difficult to get access to the system if you forget the password.

To set the root password, click **Next**.

2.4.7 Yast2: Committing your settings

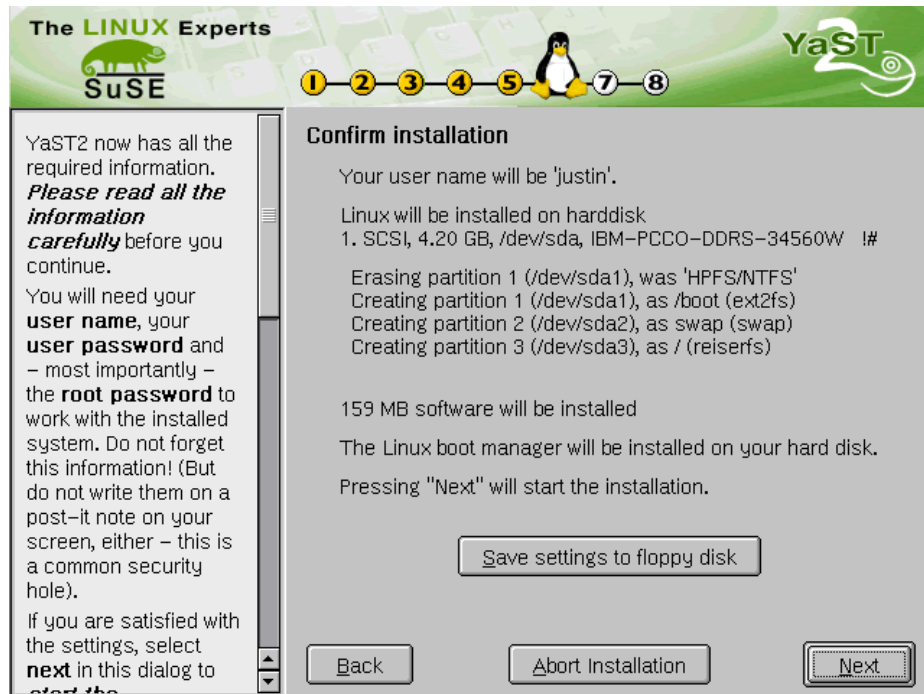


Figure 64. Summary information

This is the final window where you can still back out of committing your settings to the system. If you do not wish to start the installation, but would like to carry on at a later date, you can save the configuration to disk by clicking **Save settings to floppy disk**.

If you wish to change any settings after reviewing them, click **Back**.

To commit the changes and continue with the installation, click the **Next** button.

If you clicked the **Next** button you will see the window shown in Figure 65. Click **Yes - Install** to continue and commit your changes to the hard disk.

Click **No** to go back and change your settings or back out of the installation all together.

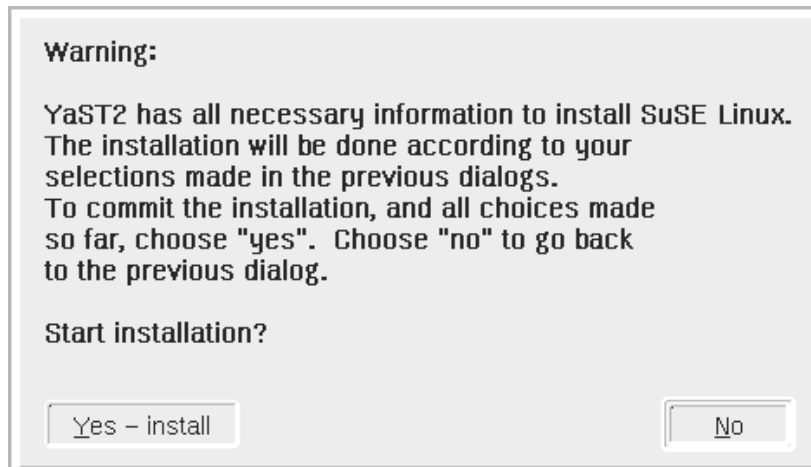


Figure 65. Selecting whether to commit the partition data

Important

Once you click **Yes-Install** your partition table will be changed to reflect your settings, and installation of packages will begin. This effectively erases all data in the partitions you have selected to change.

2.4.8 Yast2: Installing the new system

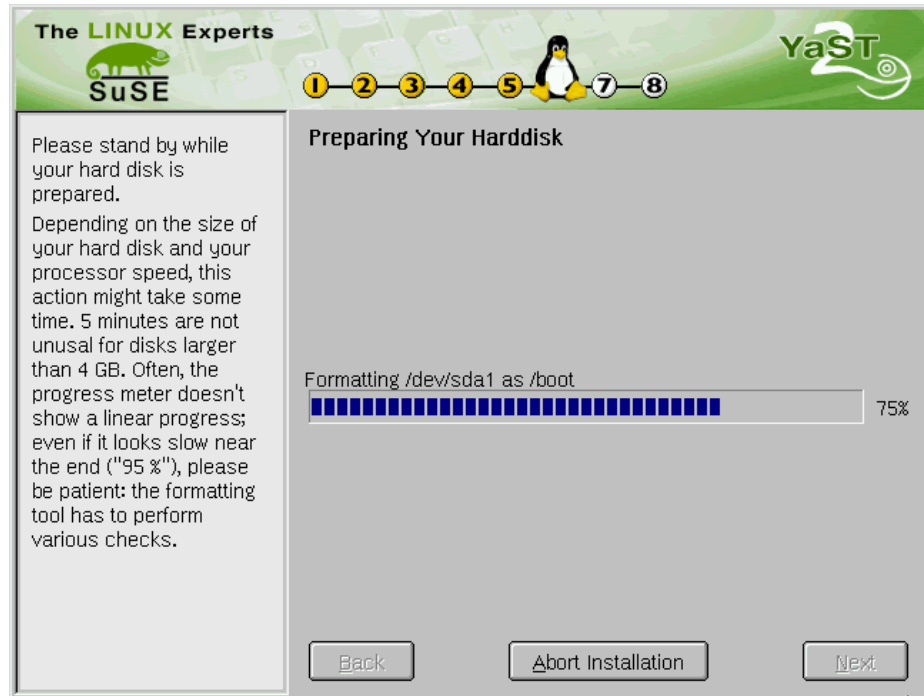


Figure 66. Formatting hard disks

Yast2 will now format your hard disks to the specifications that you entered earlier in the installation process. This may take a while, depending on the size of the partitions and hard disks in the system.

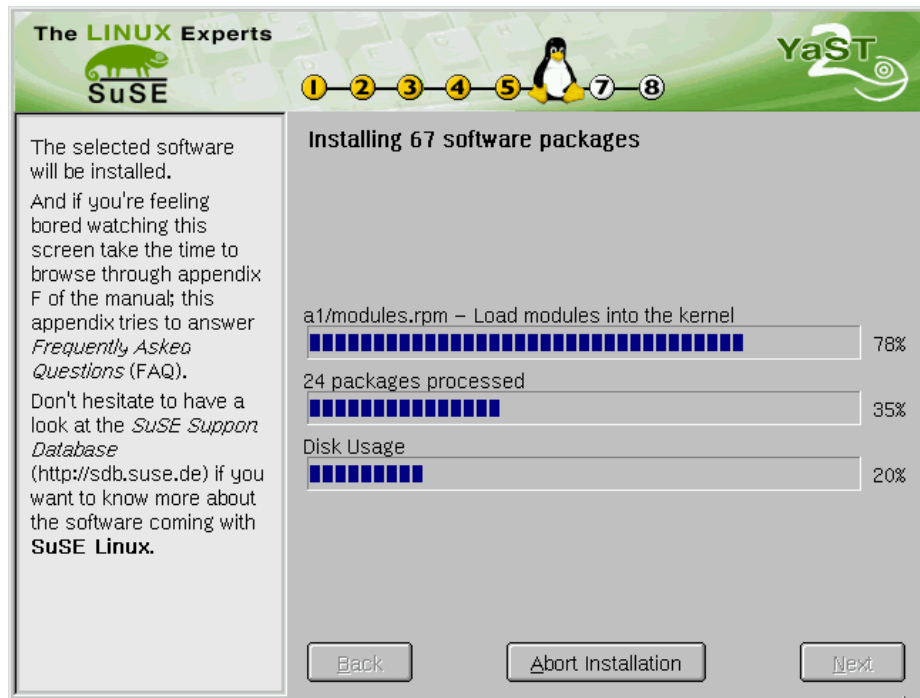


Figure 67. Installing SuSE software packages

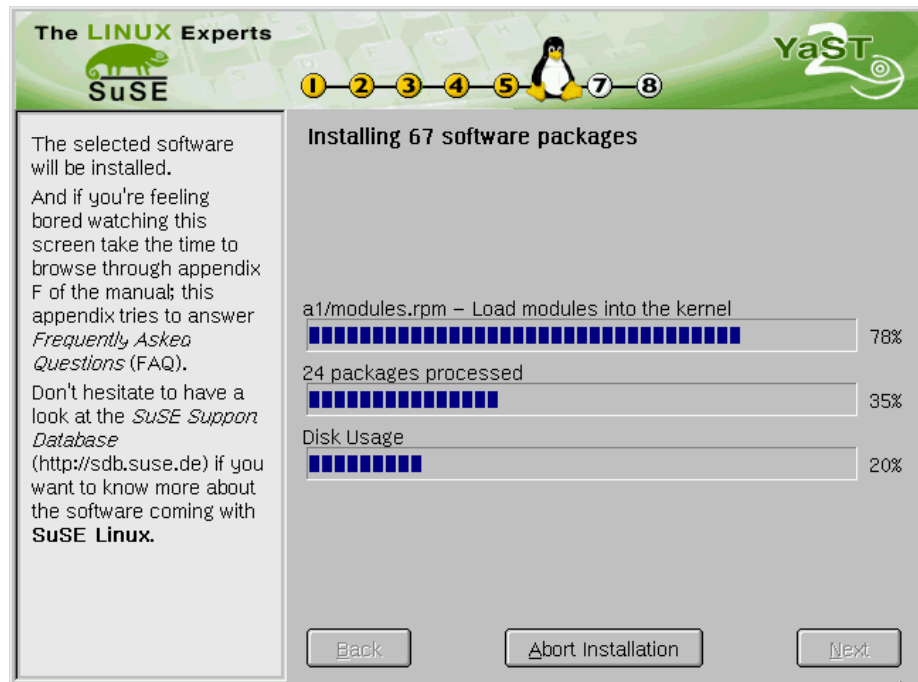


Figure 68. Installation of the base packages

Yast2 will start the installation of the packages based on the installation type you selected earlier. Depending on the amount of packages and the type of packages you installed, this may take a while.

Once the base packages are installed, Yast2 will reboot the machine to continue with the installation of the rest of the packages.

If Yast2 has detected your video settings you will proceed to the window shown in Figure 69. If your video settings could not be determined you will be presented with the window in Figure 70.

2.4.9 Yast2: Configuring your display

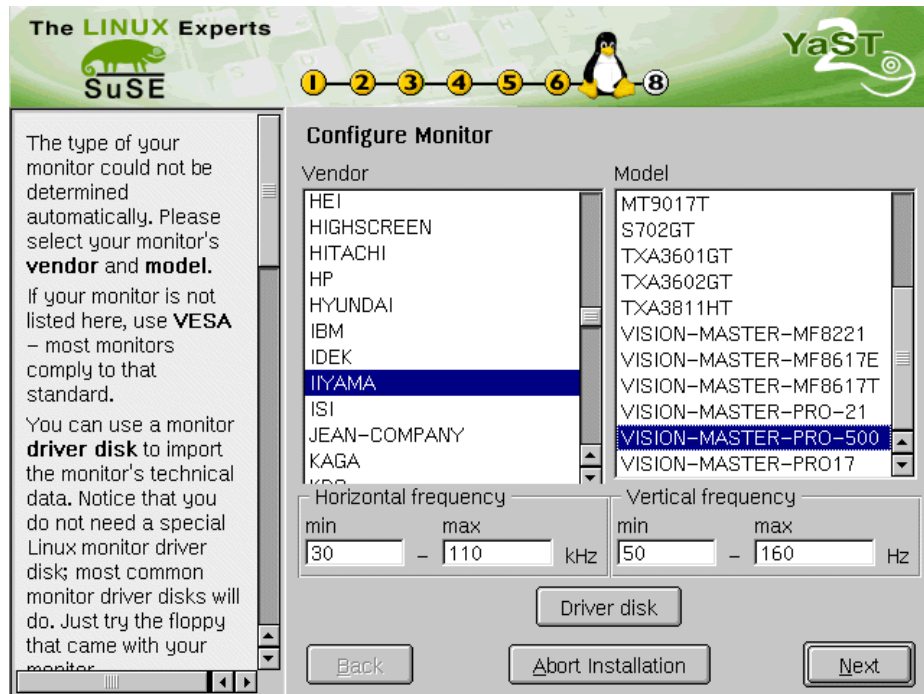


Figure 69. Configuring monitor settings

Select your monitor from the list in Figure 69. If your monitor is not listed you can choose a generic setting from the list under the LCD or VESA heading in the left hand pane.

Most monitors come with a Windows driver disk. This can be used by Yast2 to configure your monitor to the manufacturer's specifications. To use the Windows disk, click the **Driver disk** button and you can load the correct driver.

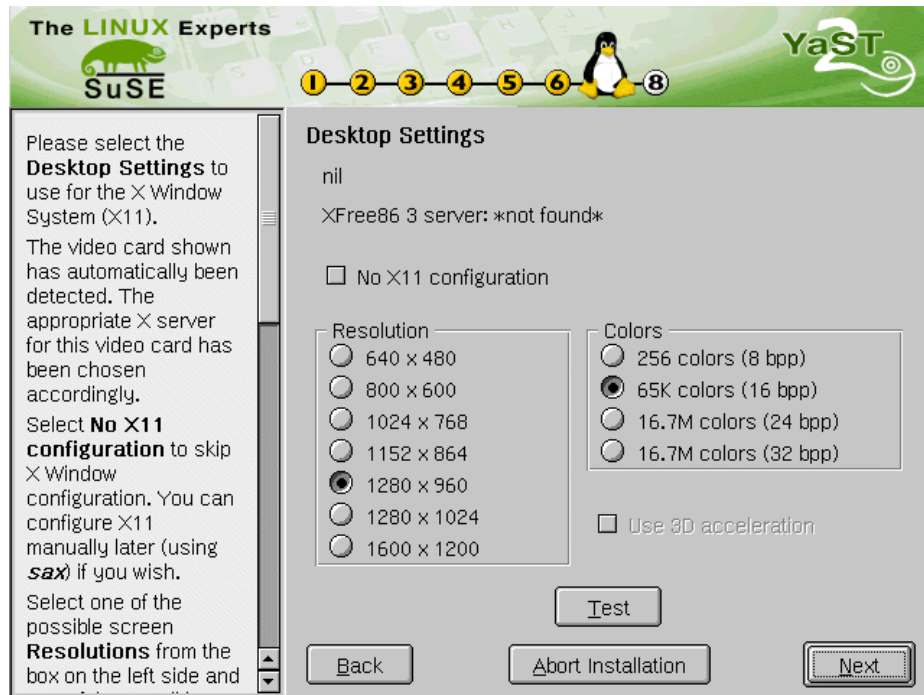


Figure 70. Configuring the video card

Once your monitor has been correctly configured, either manually or by Yast2 you are given the opportunity to configure your graphics card for use by X-Windows. If you do not wish to use X-Windows on your system, select **No X11 configuration** and click **Next**.

If you wish to configure your system for X-Windows, select the resolution and the color depth from the list. If your graphics card is supported by the SuSE 3D acceleration system, you can select **Use 3D Acceleration**.

It is always advisable to test the configuration before committing it to the system. To test the settings you have entered, click the **Test** button to see how your desktop would look.

Note

In some situations, Yast2 is unable to correctly configure your video card. We ran into this problem on a machine in the lab. If your machine locks up after clicking **Next**, restart the machine and select **No X11 configuration**. You can then manually configure X11 at a later time. The process for manually configuring X11 is covered later in this chapter.

Once your graphics system has been configured, Yast2 will check for hardware devices in your system that it can configure and set up for you.

2.4.10 Yast2: Configuring your devices

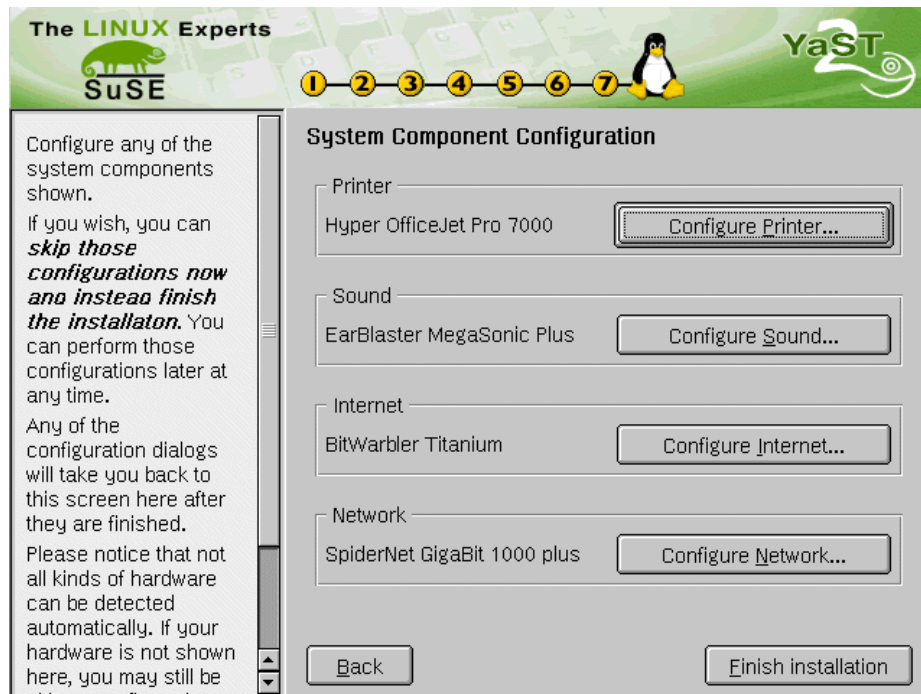


Figure 71. Configuring hardware devices

If Yast2 has found any devices under the sections shown, you are able to configure them further.

We will go through the configuration of these devices at a later stage, because the modules used for configuration of these devices are also used for administrative purposes in the Yast2 administration program.

Click **Next** to complete the SuSE installation. You can now go ahead and use some of the methods detailed in the manual and this redbook to further configure your system to your needs.

2.5 ServeRAID: Installation and configuration

In this section we will describe how to install SuSE Linux on the IBM Netfinity servers with the IBM ServeRAID controller and how to use the features of the IBM ServeRAID controller. The IBM ServeRAID controller is a high-performance RAID controller. In the current version of the Linux driver, all ServeRAID adapter versions are supported. Before you start the installation, you need to define the RAID arrays and the logical drives. The logical drives are represented to the operating system as if they were physical disk drives. For more information on RAID levels and performance issues, see Appendix A, "RAID levels" on page 361.

Stop

Before installing SuSE Linux on the IBM Netfinity server with an IBM ServeRAID controller, you need to define RAID arrays and logical drives. You can do this with ServerGuide, which comes with all IBM Netfinity servers, or with the ServeRAID DOS Configuration diskette, which is available at <http://www.pc.ibm.com/support>.

We strongly recommend that you use hot spare hard disks in your system to secure your data the best possible way.

SuSE Linux supports the ServeRAID SCSI controller. To install the operating system, follow the procedure in Chapter 2, "Installing SuSE Linux" on page 5. Your Logical disk drives defined in the RAID array will appear as SCSI drives in the installation program.

After you have installed the system, get the utilities for RAID administration from:

<http://www.developer.ibm.com/welcome/netfinity/serveraid.html>

From that site you can download the following files:

- `ipsutils.rpm`: this file contains the Linux utilities for the IBM ServeRAID SCSI adapter,
- `RaidMan.rpm`; this file contains the Linux ServeRAID Manager, which can be used to locally and remotely configure and monitor the ServeRAID controller used in Linux installation through the graphical user interface.

These files can also be found on the IBM support site:

<http://www.pc.ibm.com/support>

Here you can download diskettes with the latest firmware, drivers, utilities and ServeRAID Manager. You can also download the CD image with all the files included. The files for Linux on the CD are in directory:

`\programs\linux`

2.5.1 Installing SuSE with the latest ServeRAID driver

SuSE Linux 7.0 was released before the new IBM ServeRAID driver, and as such the support for the new 4.40 ServeRAID BIOS is not in the stock SuSE 7.0 distribution.

Note

The update of the ServeRAID driver can be found at:

<ftp://ftp.suse.com/pub/suse/i386/update/7.0/kernel/ips-4.40>

The directory contains the file `bootdisk.img`. This is an update for the installation system for ServeRAID systems.

See section 2.7.3 in the SuSE manual for instructions on how to install the floppy boot image.

The `ips-4.40` directory also contains the `ips.o` module for ServeRAID on uniprocessor systems, as well as `ips-smp.o` for multiprocessor systems.

To install on a system with the new ServeRAID BIOS, you will have to use the boot disk detailed earlier in the chapter to boot Linuxrc, and create an initial ramdisk. This is so the system will load the ServeRAID driver from the initial ramdisk to boot SuSE from the ServeRAID system. This is not as difficult as it sounds, and we will guide you through the process.

After you have created your logical drives on the system, create the boot disk from the SuSE FTP site, and also copy the `ips-smp.o` and `ips.o` to a separate floppy disk.

Insert the boot disk you created into the system and power on. Once you have gone through the Linuxrc process, select **Kernel Modules** (Figure 6 on page 11) and load the ServeRAID module.

When you are asked how you wish to install SuSE Linux, either with Yast1 or Yast2, select **Yast1**.

The installation procedure now differs as to whether you are installing on a uniprocessor system or a multiprocessor system.

Follow the steps below according to the system you are installing on:

1. Proceed with the installation with Yast1 as detailed earlier in the chapter. For a multiprocessor system, make a note of the partition you define as the root (*/*) (see Figure 15 on page 19).
2. If you are installing on a uniprocessor system, go to Step 4. For a multiprocessor system, go to Step 3.
3. Once the initial installation has finished, you will be told that the system will reboot. *Do not* take the boot disk out of the system. We need it to load the installed SuSE system.

Let the system reboot.

Once Linuxrc has loaded, go through the options as you did before, making sure to load the ServeRAID driver from the modules menu. Select **Start/Install System** from the main menu and select **Boot installed system**. Enter your root device (the partition you defined as */* in the initial installation).

The system will now boot. Proceed to Step 4.

4. Once Yast has booted the installed system, you will need to install a new initial ramdisk to allow the system to boot from the ServeRAID adapter.

Insert the disk you copied the ips modules to and at the prompt type:

```
mount /floppy
```

If you are installing on a uniprocessor system, copy the ips.o module to `/lib/modules/2.2.16/scsi/ips.o`. For a multiprocessor system, copy the ips-smp.o module to `/lib/modules/2.2.16-SMP/scsi/ips.o`

Edit the `/etc/rc.config` file. Find the line that states `INITRD_MODULES=` and add ips to the list (if there is no list, you still need to add it). If ips is already there, you do not have to add it again and you do not have to run SuSEconfig. If you have to add ips to the list, save the file and run SuSEconfig to commit the changes.

Once SuSEconfig has finished, you will need to run `mk_initrd` to create the initial ramdisk. Once the `mk_initrd` command has finished, you will need to load the `/etc/lilo.conf` to make sure a line similar to `initrd=/boot/initrd` exists. If it does not exist, enter it in the `[global]` section of the file and run `lilo` at the prompt.

The driver has been successfully installed in the initial ramdisk. You can now use the system as usual. We told you it was easy, didn't we?

2.5.2 Installing `ipsutils.rpm`

To install the `ipsutils` package you have to be logged in as `root`. After you have downloaded the `ipsutil.rpm` package, you need to install it. The `ipsutil` package is a standard Red Hat Package Manager (RPM) package.

To install the `ipsutil` package, open a terminal window, or log in at the console and execute:

```
rpm -Uhv ipsutil.rpm
```

Note

You can also copy the `ipssend` program from your ServeRAID CD-ROM to the `/usr/bin` directory with the commands:

```
mount /cdrom
```

```
cp /cdrom/programs/linux/cmdline/ipssend /usr/bin/
```

The you need to change the permissions so that you can execute the command with:

```
chmod 755 /usr/bin/ipssend
```

This assumes that your current directory is where the `ipsutil.rpm` file resides. The necessary files will be installed in the `/usr/bin` directory. To see if the utilities are working, type the following command:

```
ipssend
```

You should see an output similar to Figure 72.

```

Licensed Material - Property of IBM Corporation
IBM ServeRAID Command Line Interface v4.40.03
(C) Copyright IBM Corp. 1994, 2000. All Rights Reserved.
US Government Restricted Rights - Use, Duplication, or Disclosure
Restricted by GSA ADP Schedule Contract with IBM Corporation

Usage: IPSSEND <Command> <Param 1> ... <Param N>
Help : IPSSEND <Command> for specific help on any command.

  Command | Param 1 | Param 2 | Param 3 | Param 4 | Param 5
  -----|-----|-----|-----|-----|-----
AUTOSYNC | Controller | Logical Drive | NOPROMPT |         |
BACKUP   | Controller | Filename      | NOPROMPT |         |
DEVINFO  | Controller | Channel       | SCSI ID  |         |
DRIVEVER | Controller | Channel       | SCSI ID  |         |
ERASEEVENT | Controller | Options       |          |         |
GETCONFIG | Controller | Options       |          |         |
GETEVENT | Controller | Options       |          |         |
GETSTATUS | Controller |              |          |         |
HSREBUILD | Controller | Options       |          |         |
INIT     | Controller | Logical Drive | NOPROMPT |         |
REBUILD  | Controller | Channel       | SCSI ID  | New Channel | New SCSI ID
RESTORE  | Controller | Filename      | NOPROMPT |         |
SETSTATE | Controller | Channel       | SCSI ID  | New State  |
SYNCH   | Controller | Scope         | Scope ID |          |
UNATTENDED | Controller | Options       |          |         |
UNBLOCK  | Controller | Logical Drive |          |         |

```

Figure 72. ipssend command output

As you can see, ipssend supports quite a lot of commands for dealing with the IBM ServeRAID controller. In this section we will cover the ones that are necessary in order to use the ServerRAID controller efficiently.

2.5.3 The ipssend commands

In this section we cover the different options of the ipssend command.

2.5.3.1 The getconfig command

This command is used to get the configuration information of the IBM ServeRAID controller, the logical drives and the physical drives. The getconfig command has the following syntax:

```
ipssend getconfig <Controller> <Options>
```

The parameters are explained in Table 1.

Table 1. *getconfig* command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Options	AD for Controller Information
	LD for Logical Drive Information
	PD for Physical Device Information
	AL (default) for All Information

To get all information about the first ServeRAID controller, execute the following command:

```
ispsend getconfig 1
```

You will see a window similar to Figure 73.


```

Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Controller Information
-----
Firmware Version      : 3.73.00
Boot Block Version    : 3.00.16
BIOS Version          : 4.40.03
Controller Type       : ServeRAID-3L
Controller Slot Information : 1
Controller Configuration ID : Null Config
SCSI Channel Description : 1 parallel SCSI wide
Initiator IDs (Channel/SCSI ID) : 1/7
Maximum Physical Devices : 15
Defunct Disk Drive Count : 0
Logical Drives/Offline/Critical : 1/0/0
Rebuild Rate (Low/Medium/High) : High
Read Ahead            : Adaptive
Unattended Mode (Yes/No) : No
Part of Cluster (Yes/No) : No
Concurrent Commands Supported : 32
Configuration Update Count : 1
-----
Logical Drive Information
-----
Logical Drive Number 1
Status of Logical Drive : Okay (OKY)
Raid Level              : 5
Size (in MB)           : 52068
Write Cache Status     : Write Back (WB)
Number of Chunks       : 7
Stripe Unit Size      : 8K
Access Blocked         : No
Part of Array          : A
Part of Merge Group    : 207

Array A Stripe Order (Channel/SCSI ID) : 1,1 1,2 1,3 1,4 1,8 1,9 1,10
-----
Physical Device Information
-----
Channel #1:
Initiator at SCSI ID 7
Target on SCSI ID 0
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 0
PFA (Yes/No) : No
State : Hot Spare (HSP)
Size (in MB)/(in Sectors) : 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A27C

```

Figure 73. Executing `ipssend getconfig 1`

In this output you can see all information about the ServeRAID configuration. If you want information only about the controller itself, execute this command:

```
ipssend getconfig 1 ad
```

You will see output similar to Figure 74.

```
[root@nf3500a /root]# ipssend getconfig 1 ad
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Controller Information
-----
Firmware Version      : 3.73.00
Boot Block Version    : 3.00.16
BIOS Version          : 4.40.03
Controller Type       : ServeRAID-3L
Controller Slot Information : 1
Controller Configuration ID : Null Config
SCSI Channel Description : 1 parallel SCSI wide
Initiator IDs (Channel/SCSI ID) : 1/7
Maximum Physical Devices : 15
Defunct Disk Drive Count : 0
Logical Drives/Offline/Critical : 2/0/0
Rebuild Rate (Low/Medium/High) : High
Read Ahead           : Adaptive
Unattended Mode (Yes/No) : No
Part of Cluster (Yes/No) : No
Concurrent Commands Supported : 32
Configuration Update Count : 24
Command Completed Successfully.
```

Figure 74. Executing ipssend getconfig 1 ad

To get information about logical drives execute this command:

```
ipssend getconfig 1 ld
```

You will get output similar to Figure 75.

```
[root@nf3500a /root]# ipssend getconfig 1 ld
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Logical Drive Information
-----
Logical Drive Number 1
  Status of Logical Drive      : Okay (OKY)
  Raid Level                   : 5
  Size (in MB)                 : 2000
  Write Cache Status           : Write Through (WT)
  Number of Chunks              : 3
  Stripe Unit Size              : 8K
  Access Blocked                : No
  Part of Array                 : A
  Part of Merge Group           : 207
Logical Drive Number 2
  Status of Logical Drive      : Okay (OKY)
  Raid Level                   : 5
  Size (in MB)                 : 2000
  Write Cache Status           : Write Through (WT)
  Number of Chunks              : 3
  Stripe Unit Size              : 8K
  Access Blocked                : No
  Part of Array                 : A
  Part of Merge Group           : 207

  Array A Stripe Order (Channel/SCSI ID) : 1,1 1,2 1,3
Command Completed Successfully.
```

Figure 75. Executing ipssend getconfig 1 ld

From this output you can get all information about the logical drives:

- Drive status
- RAID Level
- Size
- Write Cache Status
- Number of Chunks
- Stripe Unit Size
- Access
- Array

To get detailed information about a physical drive, execute this command:

```
ipssend getconfig 1 pd
```

You will see output similar to Figure 76.

```
[root@nf3500a /root]# ipssend getconfig 1 pd
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Physical Device Information
-----
Channel #1:
Initiator at SCSI ID 7
Target on SCSI ID 0
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 0
PFA (Yes/No) : No
State : Ready (RDY)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A27C
Target on SCSI ID 1
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 1
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL09YSS
Target on SCSI ID 2
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 2
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A2QK
Target on SCSI ID 3
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 3
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG DMVS09D 01B0F802F9F4
```

Figure 76. Executing `ipssend getconfig 1 pd`

2.5.3.2 The `getstatus` command

This command is used to retrieve the current status of the IBM ServeRAID controller. The `getstatus` command has the following syntax:

```
ipssend getstatus <Controller>
```

Where you see the Controller parameter is the number of the controllers (from 1 to 12).

To get the status of first ServeRAID controller in your IBM Netfinity server, execute this command:

```
ipssend getstatus 1
```

You will see output similar to Figure 77.

```

[root@nf3500a /root]# ipssend getstatus 1
Found 1 IBM ServeRAID Controller(s).
Background Command Progress Status for controller 1...
  Current/Most Recent Operation : Rebuild
  Logical Drive in Progress     : 2
  Rebuild Rate                  : High
  Status                        : Successfully Completed
  Logical Drive Size (in Stripes): 128000
  Number of Remaining Stripes  : 0
  Percentage Complete           : 100.00%
Command Completed Successfully.

```

Figure 77. Executing `ipssend getstatus 1`

If the ServeRAID controller is in the middle of rebuilding a drive, you will see output similar to Figure 78.

```

[root@nf3500a /root]# ipssend getstatus 1
Found 1 IBM ServeRAID Controller(s).
Background Command Progress Status for controller 1...
  Current/Most Recent Operation : Rebuild
  Logical Drive in Progress     : 1
  Rebuild Rate                  : High
  Status                        : In Progress
  Logical Drive Size (in Stripes): 128000
  Number of Remaining Stripes  : 126473
  Percentage Complete           : 1.19%
Command Completed Successfully.

```

Figure 78. Executing `ipssend getstatus 1` during rebuilding of a drive

2.5.3.3 The `devinfo` command

This command is used to retrieve the current status of the devices connected to the IBM ServeRAID controller. The `devinfo` command has the following syntax:

```
ipssend devinfo <Controller> <Channel> <SCSI ID>
```

The parameters are explained in Table 2.

Table 2. devinfo command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Channel	Channel of Device (1 to 3)
SCSI ID	SCSI ID of Device (0 to 15)

To get the status of a device with SCSI ID 1 on channel 1 on the first ServeRAID controller in your IBM Netfinity server, execute the command:

```
ipssend devinfo 1 1 1
```

You will see output similar to Figure 79.

```
[root@nf3500a /root]# ipssend devinfo 1 1 1
Found 1 IBM ServeRAID Controller(s).
Device Information has been initiated for controller 1...
  Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
  Channel          : 1
  SCSI ID         : 1
  PFA (Yes/No)    : No
  State           : Online (ONL)
  Size (in MB)/(in Sectors): 8678/17773888
  Device ID      : IBM-PSG ST39175L04303AL09YSS
Command Completed Successfully.
```

Figure 79. Executing ipssend devinfo 1 1 1

If the ServeRAID controller is in the middle of rebuilding a drive, you will see output similar to Figure 80.

```
[root@nf3500a /root]# ipssend devinfo 1 1 2
Found 1 IBM ServeRAID Controller(s).
Device Information has been initiated for controller 1...
  Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
  Channel          : 1
  SCSI ID         : 2
  PFA (Yes/No)    : No
  State           : Rebuild (RBL)
  Size (in MB)/(in Sectors): 8678/17773888
  Device ID      : IBM-PSG ST39175L04303AL0A2QK
Command Completed Successfully.
```

```

[root@nf3500a /root]# ipssend devinfo 1 1 2
Found 1 IBM ServeRAID Controller(s).
Device Information has been initiated for controller 1...
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
Channel          : 1
SCSI ID          : 2
PFA (Yes/No)    : No
State            : Rebuild (RBL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID       : IBM-PSG ST39175L04303AL0A2QK
Command Completed Successfully.

```

Figure 80. Executing `ipssend devinfo 1 1 2` during rebuilding of a drive

2.5.3.4 The `hsrebuild` command

This command is used for setting the state of the Hot Swap Rebuild option. The `hsrebuild` command has the following syntax:

```
ipssend hsrebuild <Controller> <Options>
```

The parameters are explained in Table 3.

Table 3. `hsrebuild` command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Options	ON: enable Hot Swap Rebuild
	?: Display status of Hot Swap Rebuild feature

With this command you can retrieve or set the Hot Swap Rebuild feature. If the Hot Swap Rebuild feature is ON, it means that if one drive in the RAID array fails, rebuilding of this drive will start automatically when you replace the failed drive with a new one. This can improve the safety of your data.

Note

The Hot Swap Rebuild feature should not be confused with a hot spare drive. A hot spare drive means that a drive is in a waiting state as long as the RAID array is in an Okay state. Once the RAID array becomes in a Critical state, the hot spare drive is enabled and the data from the defunct drive automatically gets rebuilt onto the hot spare drive, disregarding the Hot Swap Rebuild setting.

To retrieve the information about the Hot Swap Rebuild status on the first ServeRAID controller, execute this command:

```
ipssend hsrebuild 1 ?
```

You will see output similar to Figure 81.

```
[root@nf3500a /root]# ipssend hsrebuild 1 ?  
Found 1 IBM ServeRAID Controller(s).  
Set Hot Swap Rebuild has been initiated for controller 1...  
Hot Swap Rebuild is On for controller 1.
```

Figure 81. Executing `ipssend hsrebuild 1 ?`

To enable the Hot Swap Rebuild option, execute this command:

```
ipssend hsrebuild 1 on
```

You will see output similar to Figure 82.

```
[root@nf3500a /root]# ipssend hsrebuild 1 on  
Found 1 IBM ServeRAID Controller(s).  
Set Hot Swap Rebuild has been initiated for controller 1...  
Hot Swap Rebuild is already On for controller 1.
```

Figure 82. Executing `ipssend hsrebuild 1 on`

2.5.3.5 The `setstate` command

With the `setstate` command you redefine the state of a physical device from the current state to the designated state. The `setstate` command has the following syntax:

```
ipssend setstate <Controller> <Channel> <SCSI ID> <New State>
```

The parameters are explained in Table 4.

Table 4. *setstate* command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Channel	Channel of device (1 to 3)
SCSI ID	SCSI ID of device (0 to 15)
New State	EMP (Empty) RDY (Ready) HSP (Hot Spare) SHS (Standby Hot Spare) DDD (Defunct Disk Drive) DHS (Defunct Hot Spare) RBL (Rebuild) SBY (Standby) ONL (Online)

Stop

Extreme caution must be taken when executing this command! Redefining a defunct (DDD) device to online (ONL) without going through a rebuild is extremely dangerous.

Before changing the state of a physical device, you can check the current status with this command:

```
ipssend getconfig 1 pd
```

With this command you will see all physical devices, except empty ones, on the first IBM ServeRAID controller. For example if you want to set the state of device on the first ServeRAID controller, channel 1 and SCSI ID 0 to RDY - Ready, execute this command:

```
ipssend setstate 1 1 0 rdy
```

You will see output similar to Figure 83.

```
[root@nf3500a /root]# ipssend setstate 1 1 0 rdy
Found 1 IBM ServeRAID Controller(s).
Set Device State has been initiated for Controller 1...
Command Completed Successfully.
```

Figure 83. Executing *ipssend setstate 1 1 0 rdy*

You can verify the change of the device state by executing this command:

```
ipssend getconfig 1 pd
```

2.5.3.6 The `synch` command

This command is used to synchronize the parity information on redundant logical drives. If the parity information is inconsistent, it will automatically be repaired. The `synch` command has the following syntax:

```
ipssend synch <Controller> <Scope> <Scope ID>
```

The parameters are explained in Table 5.

Table 5. `synch` command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Scope	Drive for a single logical drive
Scope ID	Number of logical drive (1 to 8)

Note

We recommend that you use this command on a weekly basis.

2.5.3.7 The `unattended` command

This command is used to alter the unattended mode of the ServeRAID controller. The `unattended` command has the following syntax:

```
ipssend unattended <Controller> <Options>
```

The parameters are explained in Table 6.

Table 6. `unattended` command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Options	ON: enable unattended mode
	OFF: disable unattended mode
	?: display status of unattended mode feature

If you want to see the current status of the first ServeRAID controller, execute this command:

```
ipssend unattended 1 ?
```

You will see output similar to Figure 84.

```
[root@nf3500a /root]# ipssend unattended 1 ?
Found 1 IBM ServeRAID Controller(s).
Set Unattended Mode has been initiated for controller 1...
Unattended Mode is set Off.
```

Figure 84. Executing `ipssend unattended 1 ?`

If you want to set the unattended mode to ON, execute this command:

```
ipssend unattended 1 on
```

You will see output similar to Figure 85.

```
[root@nf3500a /root]# ipssend unattended 1 on
Found 1 IBM ServeRAID Controller(s).
Set Unattended Mode has been initiated for controller 1...
Command Completed Successfully.
```

Figure 85. Executing `ipssend unattended 1 on`

2.5.3.8 The rebuild command

This command starts a rebuild to the designated drive. The `rebuild` command has the following syntax:

```
ipssend rebuild <Controller> <Channel> <SCSI ID> <New Channel> <New SCSI ID>
```

The parameters are explained in Table 7.

Table 7. `rebuild` command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Channel	Channel of defunct drive (1 to 3)
SCSI ID	SCSI ID of defunct drive (0 to 15)
New Channel	Channel of new drive (1 to 3)
New SCSI ID	SCSI ID of new drive (0 to 15)

This operation is valid for disk arrays containing one or more logical drives in a Critical (CRT) state. For example, if you want to rebuild a defunct drive on

SCSI ID 2 on channel 1 on the first ServerRAID controller to a new drive on SCSI ID 0 on the same channel, you will execute this command:

```
ipssend rebuild 1 1 2 1 0
```

You will see output similar to Figure 86.

```
[root@nf3500a /root]# ipssend`rebuild 1 1 2 1 0
Found 1 IBM ServeRAID Controller(s).
Rebuild Drive has been initiated for controller 1...
Rebuilding Logical Drive #1:
.....10% Done
.....20% Done
.....30% Done
.....40% Done
.....50% Done
.....60% Done
.....70% Done
.....80% Done
.....90% Done
..... Done Logical Drive #1
Rebuilding Logical Drive #2:
.....10% Done
.....20% Done
```

Figure 86. Executing `ipssend rebuild 1 1 2 1 0`

2.5.4 Replacing a defunct drive

When a physical drive in a RAID array becomes defunct you will see a light signal on the drive. You can simulate a defunct drive by executing the following command:

```
ipssend setstate 1 1 3 ddd
```

In this case we are simulating that the drive with SCSI ID 3 on channel 1 on the first ServeRAID controller is defunct. The following steps should be taken to replace the defunct drive:

1. Physically replace the defunct drive with a good drive.
2. The IBM ServeRAID controller will start rebuilding the drive automatically.

Note

Automatically rebuilding will work only on ServeRAID II and III. And Enable Hot Spare Rebuild must be set to Enabled!

You can check the progress of rebuilding the logical drives on the first IBM ServeRAID controller with this command:

```
ipssend getstatus 1
```

You will see output similar to Figure 78 on page 79.

If the rebuild is not completed successfully, you will see output similar to Figure 87.

```
[root@nf3500a /root]# ipssend getstatus 1
Found 1 IBM ServeRAID Controller(s).
Background Command Progress Status for controller 1...
  Current/Most Recent Operation : Rebuild
  Logical Drive in Progress      : 1
  Rebuild Rate                   : High
  Status                         : Drive Failed
  Channel Number is             : 1
  SCSI ID Number is             : 0
  Logical Drive Size (in Stripes): 128000
  Number of Remaining Stripes   : 89562
  Percentage Complete           : 30.03%
Command Completed Successfully.
```

Figure 87. Failed rebuild

2.5.5 Replacing a defunct drive with disabled Hot Spare Rebuild

When you have disabled the Hot Spare Rebuild function in the IBM ServeRAID controller configuration, the following steps should be taken to replace the defunct drive. In our example, the drive with SCSI ID 1 on channel 1 on the first ServeRAID controller is defunct.

1. Physically replace the defunct drive with a working one.
2. Execute the following command to start rebuilding the drive:

```
ipssend setstate 1 1 3 rbl
```

You will see output similar to this:

```
[root@nf3500a /root]# ipssend setstate 1 1 3 rbl
Found 1 IBM ServeRAID Controller(s).
Set Device State has been initiated for Controller 1...
Command Completed Successfully.
```

Figure 88. Forced rebuild of the defunct drive

You can check the progress of rebuilding the logical drives on the first IBM ServeRAID controller with this command:

```
ipssend getstatus 1
```

You will see output similar to Figure 78 on page 79.

2.5.6 Replacing a defunct drive with a hot spare drive installed

When you have configured the hot spare drive in your IBM ServeRAID configuration, the defunct physical drive is automatically rebuilt to the hot spare drive. Follow these steps to replace the defunct physical drive and set it as a hot spare drive:

1. You find out that there is a defunct physical drive in your RAID array on the first ServeRAID controller. In our example the physical drive on SCSI ID 2 on channel 1 was defined as a hot spare drive. You can check this by executing the command:

```
ipssend getconfig 1 pd
```

You will see output similar to Figure 89.

```
[root@nf3500a /root]# ipssend getconfig 1 pd
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Physical Device Information
-----
Channel #1:
Initiator at SCSI ID 7
Target on SCSI ID 0
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 0
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A27C
Target on SCSI ID 1
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 1
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL09YSS
Target on SCSI ID 2
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 2
PFA (Yes/No) : No
State : Rebuild (RBL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A2QK
Target on SCSI ID 3
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 3
PFA (Yes/No) : No
State : Defunct Hot Spare (DHS)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG DNES-309SAHRAJLJ6230
Target on SCSI ID 15
Device is a 16 bit, Fast SCSI, tag queuing Processor Device
SCSI ID : 15
PFA (Yes/No) : No
State : Standby (SBY)
Size (in MB)/(in Sectors): 0/0
Device ID : IBM EXP200 10D792063452
Command Completed Successfully.
```

Figure 89. After failing the drive in RAID array

As you can see, the hot spare drive is already rebuilding and the defunct drive is in Defunct Hot Spare (DHS) state.

- 2. Remove the defunct drive from the server. In our example this is the drive with SCSI ID 3 on channel 1.
- 3. Set the state of the drive to Empty (EMP) with the command:

```
ipssend setstate 1 1 3 emp
```

You will see output similar to Figure 90.

```
[root@nf3500a /root]# ipssend setstate 1 1 3 emp
Found 1 IBM ServeRAID Controller(s).
Set Device State has been initiated for Controller 1...
Command Completed Successfully.
```

Figure 90. Setting the DHS to EMP

You can check the result of this operation by executing this command:

```
ipssend getconfig 1 pd
```

You will see output similar to Figure 91.

```
[root@nf3500a /root]# ipssend getconfig 1 pd
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Physical Device Information
-----
Channel #1:
  Initiator at SCSI ID 7
  Target on SCSI ID 0
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID          : 0
    PFA (Yes/No)     : No
    State            : Online (ONL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID        : IBM-PSG ST39175L04303AL0A27C
  Target on SCSI ID 1
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID          : 1
    PFA (Yes/No)     : No
    State            : Online (ONL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID        : IBM-PSG ST39175L04303AL09YSS
  Target on SCSI ID 2
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID          : 2
    PFA (Yes/No)     : No
    State            : Rebuild (RBL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID        : IBM-PSG ST39175L04303AL0A2QK
  Target on SCSI ID 15
    Device is a 16 bit, Fast SCSI, tag queuing Processor Device
    SCSI ID          : 15
    PFA (Yes/No)     : No
    State            : Standby (SEBY)
    Size (in MB)/(in Sectors): 0/0
    Device ID        : IBM      EXP200  10D792063452
Command Completed Successfully.
```

Figure 91. After removing defunct drive

As you can see, there is no entry for the defunct drive anymore.

4. Insert a new drive into the server. In our example this will be the same location as the defunct drive.
5. Set the state of that drive to Ready (RDY) with this command:

```
ipssend setstate 1 1 3 rdy
```

You will see output similar to Figure 92.

```
[root@nf3500a /root]# ipssend setstate 1 1 3 rdy
Found 1 IBM ServeRAID Controller(s).
Set Device State has been initiated for Controller 1...
Command Completed Successfully.
```

Figure 92. Setting the new drive state to RDY

With setting the state to Ready (RDY) the drive is started.

Note

All new drives must first be set to ready (RDY).

You can check the result of this operation by executing this command:

```
ipssend getconfig 1 pd
```

You will see output similar to Figure 93.

```
[root@nf3500a /root]# ipssend getconfig 1 pd
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Physical Device Information
-----
Channel #1:
  Initiator at SCSI ID 7
  Target on SCSI ID 0
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID : 0
    PFA (Yes/No) : No
    State : Online (ONL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID : IBM-PSG ST39175L04303AL0A27C
  Target on SCSI ID 1
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID : 1
    PFA (Yes/No) : No
    State : Online (ONL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID : IBM-PSG ST39175L04303AL09YSS
  Target on SCSI ID 2
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID : 2
    PFA (Yes/No) : No
    State : Rebuild (RBL)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID : IBM-PSG ST39175L04303AL0A2QK
  Target on SCSI ID 3
    Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
    SCSI ID : 3
    PFA (Yes/No) : No
    State : Ready (RDY)
    Size (in MB)/(in Sectors): 8678/17773888
    Device ID : IBM-PSG DNES-309SAHRAJLJ6230
  Target on SCSI ID 15
    Device is a 16 bit, Fast SCSI, tag queuing Processor Device
    SCSI ID : 15
    PFA (Yes/No) : No
    State : Standby (SBY)
    Size (in MB)/(in Sectors): 0/0
    Device ID : IBM EXP200 10D792063452
Command Completed Successfully.
```

Figure 93. After setting the state to RDY

As you can see, the new drive appears as a Ready (RDY) device, in our example under SCSI ID 3 on channel 1.

6. Change the state of the new drive to Hot Spare (HSP) with this command:

```
ipssend setstate 1 1 3 hsp
```

You will see output similar to Figure 94.

```
[root@nf3500a /root]# ipssend setstate 1 1 3 hsp  
Found 1 IBM ServerRAID Controller(s).  
Set Device State has been initiated for Controller 1...  
Command Completed Successfully.
```

Figure 94. Changing the state to HSP

You can check the result of this operation by executing this command:

```
ipssend getconfig 1 pd
```

You will see output similar to Figure 95.

```
[root@nf3500a /root]# ipssend getconfig 1 pd
Found 1 IBM ServeRAID Controller(s).
Read Configuration has been initiated for Controller 1...
-----
Physical Device Information
-----
Channel #1:
Initiator at SCSI ID 7
Target on SCSI ID 0
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 0
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A27C
Target on SCSI ID 1
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 1
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL09YSS
Target on SCSI ID 2
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 2
PFA (Yes/No) : No
State : Online (ONL)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG ST39175L04303AL0A20K
Target on SCSI ID 3
Device is a 16 bit, Fast SCSI, tag queuing Hard Disk
SCSI ID : 3
PFA (Yes/No) : No
State : Hot Spare (HSP)
Size (in MB)/(in Sectors): 8678/17773888
Device ID : IBM-PSG DNE5-309SAHRAJLJ6230
Target on SCSI ID 15
Device is a 16 bit, Fast SCSI, tag queuing Processor Device
SCSI ID : 15
PFA (Yes/No) : No
State : Standby (SBY)
Size (in MB)/(in Sectors): 0/0
Device ID : IBM EXP200 10D792063452
Command Completed Successfully.
```

Figure 95. After setting the state to HSP

Congratulations! You have just installed a brand new the new hot spare drive and it is ready to use.

2.5.7 Using the ServeRAID Manager utility

The ServeRAID Manager for Linux allows you to manage your ServeRAID controller from Linux, without the need for the not-so-stable Windows control workstation. It is a Java based tool and has the same functionality across all supported platforms. With ServeRAID Manager for Linux you can manage the ServeRAID controller locally or remotely. That means that you can install it on the server with the ServeRAID controller and manage the controller in the

server, or you can install it on a separate Linux box and manage the ServeRAID remotely.

Note

For remote management you also need to install the ServeRAID Manager on the server with the ServeRAID controller, because the agent needed for remote management is included in the package. Also the server and management station have to be connected with TCP/IP.

After you get the file `RaidMan-4.40-03.i386.rpm` from the Web or from the CD, install it with the command:

```
rpm -ihv RaidMan-4.40-03.i386.rpm
```

Note

Before using this version of the ServeRAID Manager software, the BIOS/firmware and the driver for the controller must be on the same level.

During the installation you have the option to enable the background agent which is then used for remote management. If you plan to manage the ServeRAID adapter remotely, you should answer yes. If you answer yes the installation program will add the following line to the `/etc/inittab` file:

```
nfra:123456:once:/usr/RaidMan/RaidAgnt.sh #RaidMan
```

This will start the agent in every runlevel. The installation program also starts the agent right after installation, so you do not need to reboot to start using remote management, as in some other operating systems.

The ServeRAID Manager is installed in the `/usr/RaidMan` directory.

The installation program also installs the necessary Java runtime. This Java runtime will not interfere with an already installed Java environment.

To start the ServeRAID Manager, simply execute the following command in the X Windows environment:

```
/usr/RaidMan/RaidMan.sh
```

Note

To use the ServeRAID Manager you have to have a working X Windows setup.

During the program startup you will see a window similar to Figure 96.



Figure 96. ServeRAID Manager startup

After the program is started, you will see a window similar to Figure 98.

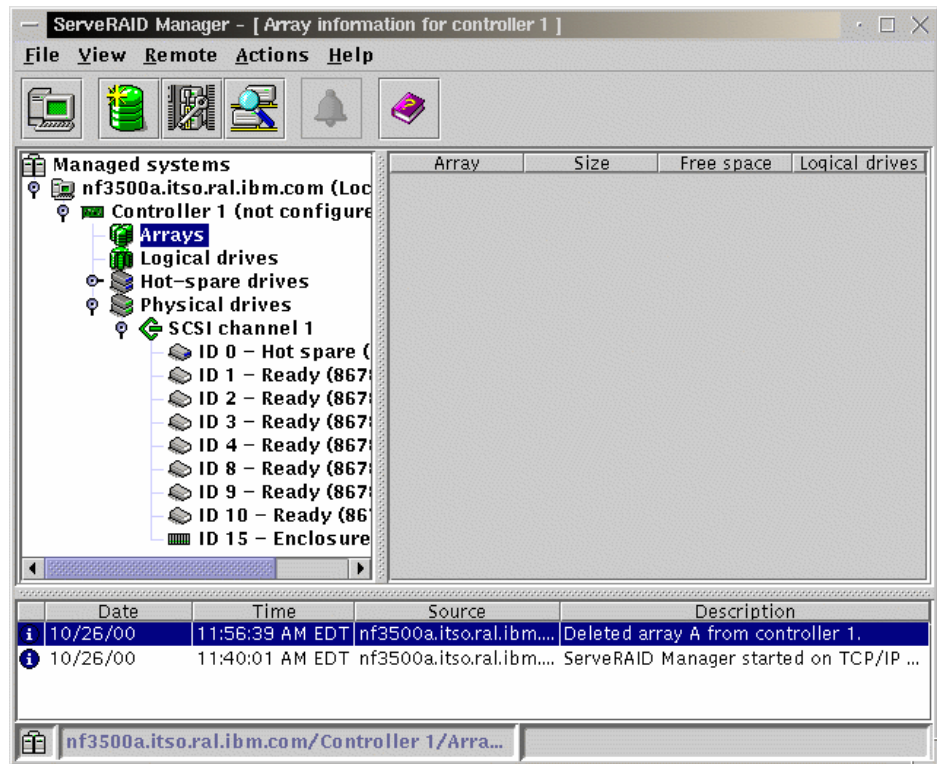


Figure 97. ServeRAID Manager

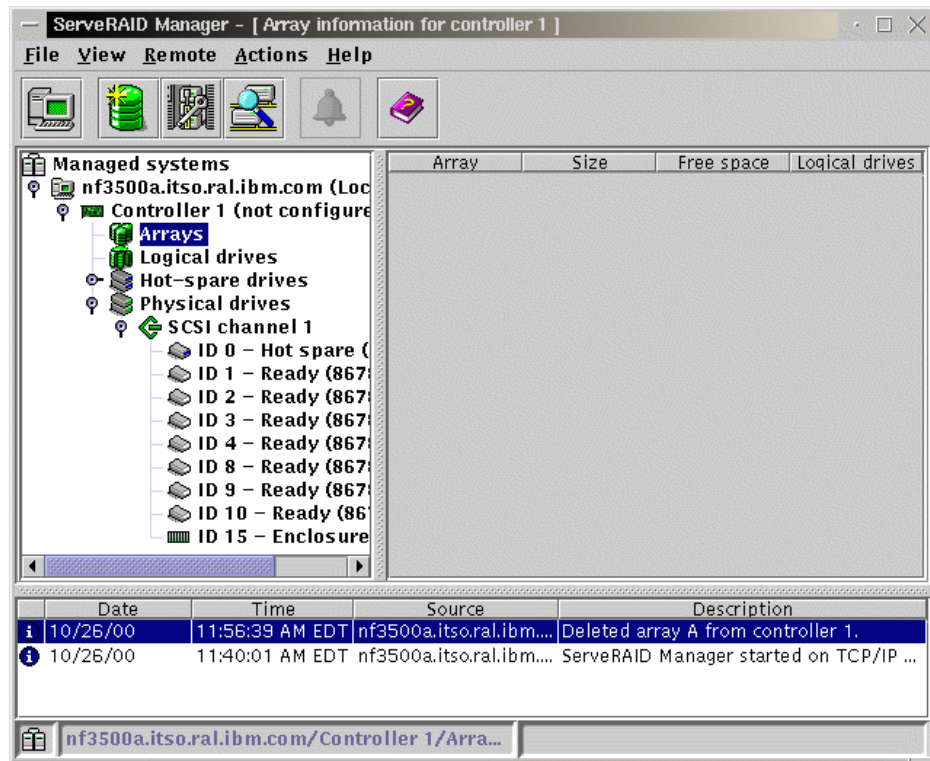


Figure 98. ServeRAID Manager

As you can see the window is divided into several areas:

- Menu bar - in the menus you can access all the functions available
- Icon bar - icons offers you shortcuts to the most often used functions
- Tree window - here you can see all the systems managed by the ServeRAID controller
- Info window - here you can see information about arrays and logical drives
- Event log - all the events are displayed here.

Note

Instructions on how to use the ServeRAID Manager functions can be found in the online help.

2.5.8 Remote management of the ServeRAID adapter

Your server with an installed ServeRAID controller can also be managed remotely. For this you need to do the following:

1. Install ServeRAID Manager on the server with installed ServeRAID adapter, as we described in 2.5.7 “Using the ServeRAID Manager utility” on page 94. Do not forget to enable the ServeRAID Manager agent running as service at boot time.
2. Install ServeRAID Manager on the Linux workstation with the TCP/IP connection to the server you would like to manage.

Before you can start using the remote management function of the ServeRAID Manager for Linux, you need to change the security setting on the server you would like to manage. If the server has a properly configured X Windows environment, you can locally start the ServeRAID Manager and update the security information. If you do not use the X Windows on your server, you need to follow these steps to enable you to access your server. By default the security is enabled after the installation of the ServeRAID Manager.

1. On the Linux workstation with the installed ServeRAID Manager, start the manager with the command:

```
/usr/RaidMan/RaidMan.sh
```

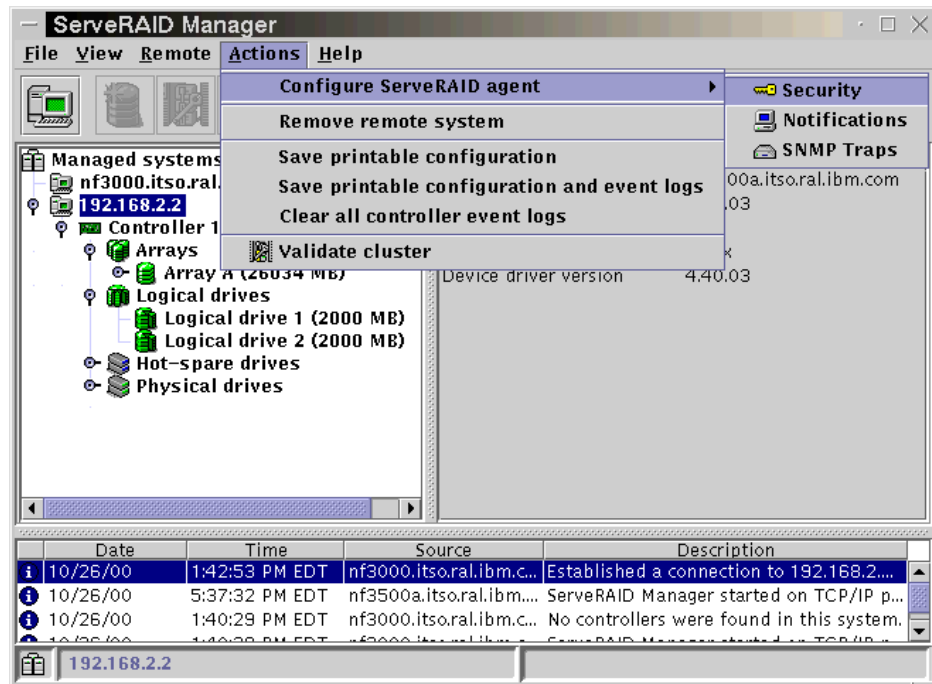



Figure 99. Starting security configuration

- From the **Actions** menu (Figure 99) select **Configure ServeRAID agent->Security**. You will see a window similar to Figure 100.

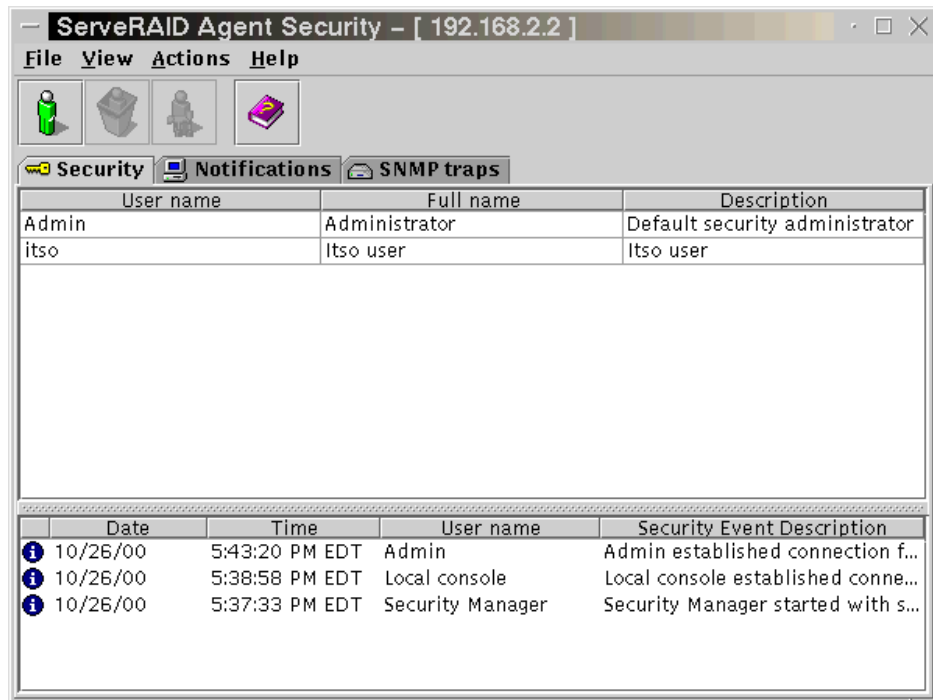


Figure 100. Security window

3. Double-click **Admin** (this is a built-in user, which cannot be removed) and you will see the window similar to Figure 101.

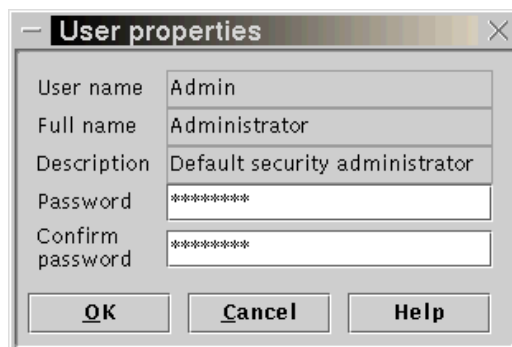


Figure 101. Changing the user

4. Type in the required password and click **OK**.
5. Close the ServeRAID Manager.

6. Install the ServeRAID Manager on the server to be managed. Copy the file `/usr/RaidMan/RaidSLst.ser` from the workstation to the server.

Note

The directory on the server has to be the same as on the workstation.

7. The command:

```
ps ax | grep jre*
```

finds all jre processes and deletes them. You also need to delete the ServeRAID Manager agent. You find the process ID with the command:

```
ps ax | grep RaidAgnt*
```

8. Using this command from the command prompt, you can restart the ServeRAID agent:

```
/usr/RaidMan/RaidMan.sh &
```

Congratulations! Your server is now ready for remote ServeRAID management.

You can connect to the remote server from the management workstation by selecting **Remote ->Add remote system** as shown in Figure 102.

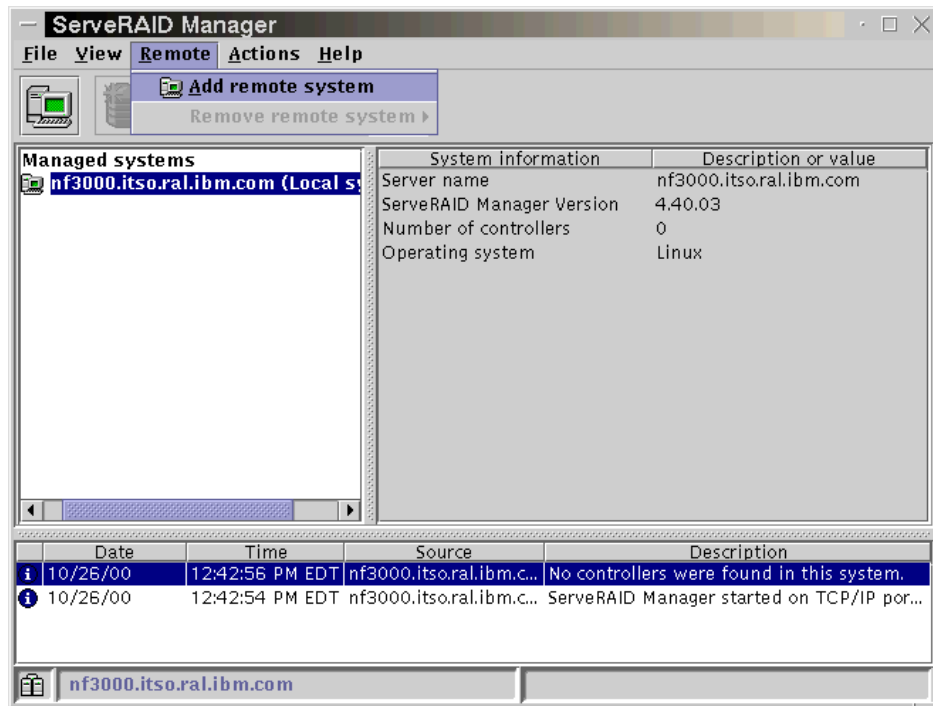


Figure 102. Accessing the remote server

You will see a window similar to Figure 103.




Figure 103. Remote system

Type in the necessary data and click **Connect**. After the system is connected, you will see a window similar to Figure 104.

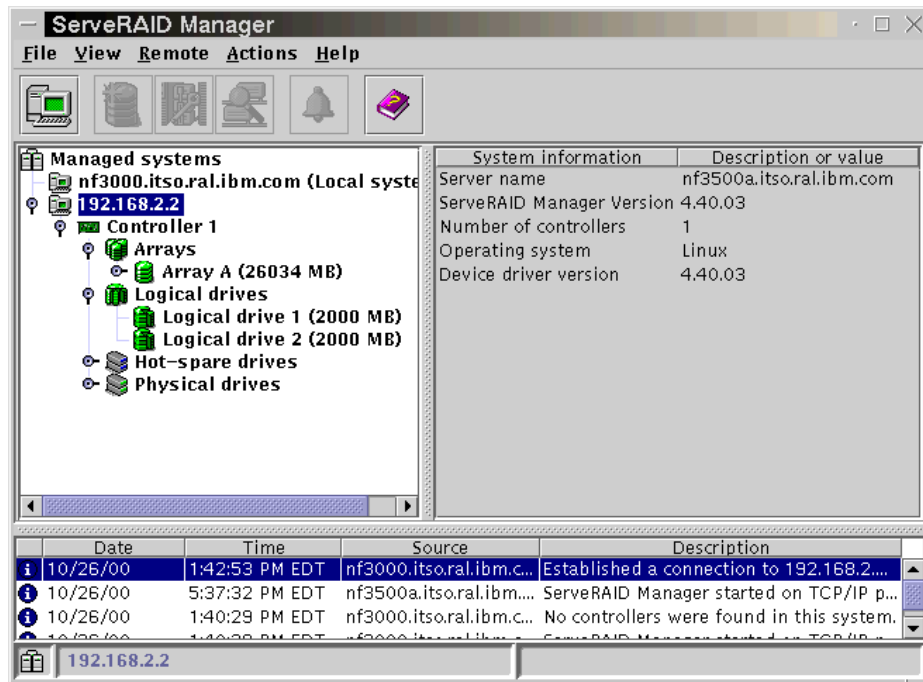


Figure 104. Remote system after connection

Now you can start managing the remote ServeRAID adapter.

2.6 XFree86 configuration

After the initial installation, the system will boot up only on the text console. While this is fine if you want to use Linux only as a server operating system, many people prefer a MS Windows-based user interface. If you want to use a graphical desktop environment like KDE or GNU Network Object Model Environment (GNOME), you first have to configure the X-Windows system to fit your configuration.

To configure XFree86, we recommend you use SaX, SuSE's advanced X configuration tool. For a more detailed documentation of SaX, see section 8.4, "Configuration using SaX" in the SuSE Linux manual. If SaX fails for some reason, you can still use XF86Setup or xf86config as a fallback solution. Both belong to the XFree86 tool collection and can also be used to configure XFree86. However, they are not as user friendly as SaX and you may need some more experience with XFree86 to use them.

For information about configuring XFree 4, please see section 8.3 “Configuration with Sax2” in the SuSE Linux manual.

To start SaX, type `sax` on the command line after you have logged in as user `root`. If you already know which X server (the “driver”) is the correct one for your video card, you can also use `sax -s <servername>`, for example `sax -s svga`.

After SaX finishes loading its configuration data, you will be presented with the mouse configuration dialog shown in Figure 105.

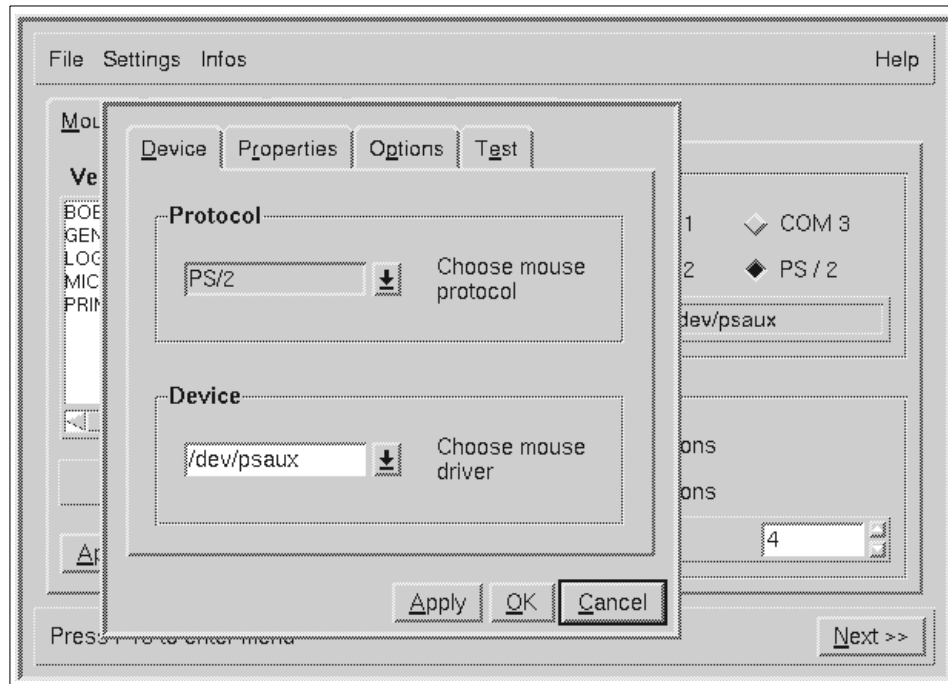


Figure 105. SaX: mouse configuration window

Because SaX is a graphical configuration tool, you first need to configure your mouse to be able to operate SaX more conveniently. If you configured your mouse during the initial system installation, you should already be able to move the mouse. If not, you have to use the keyboard by pressing the Tab key to move between the different input fields.

If your mouse is moving fine, click the folder named **Test** to test your mouse.

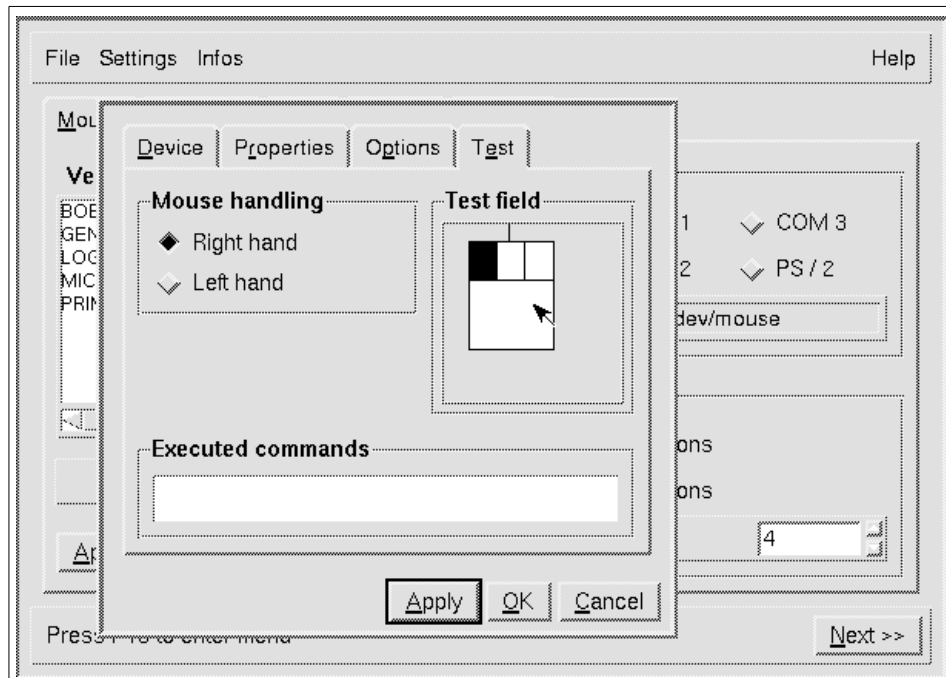


Figure 106. SaX: mouse test window

If your mouse has only two buttons, you can emulate the third (middle) mouse button by pressing the left and right button at the same time. To activate this emulation mode, select the **Options** folder and check **Emulate 3 buttons**.

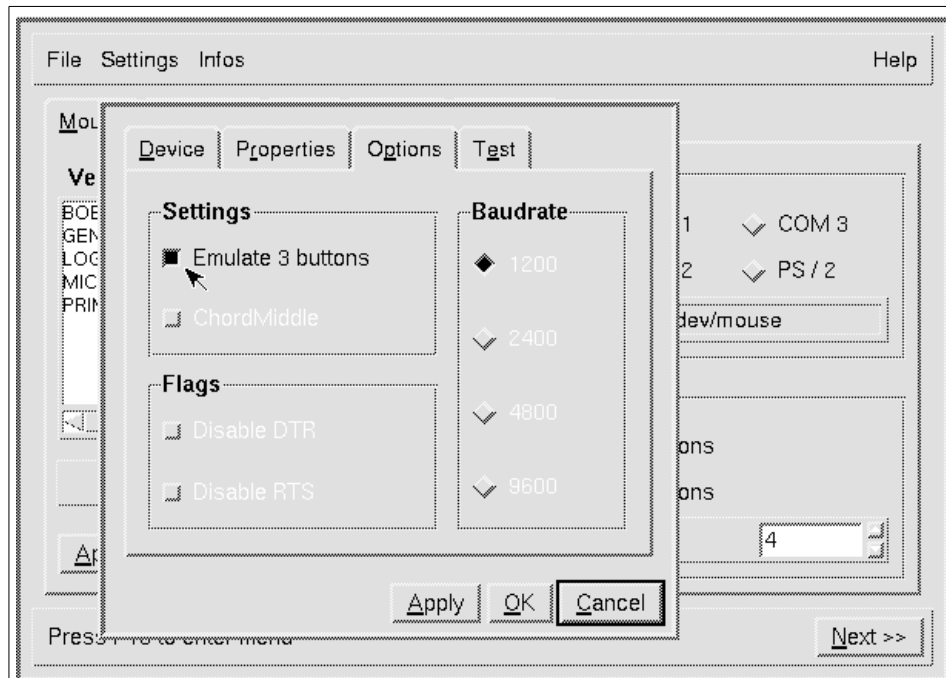


Figure 107. SaX: mouse options

Click **Apply** to apply the change. If your mouse is working fine, click **OK** to close the mouse configuration dialog. Click **Next >>** in the bottom right corner to continue to the keyboard configuration window.

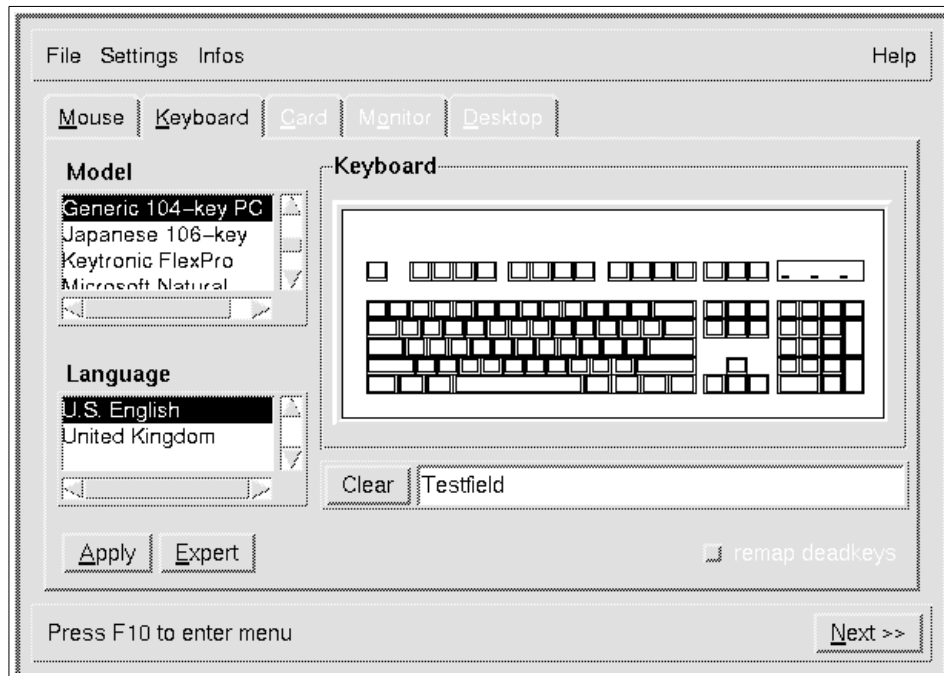


Figure 108. SaX: Keyboard configuration

By default, SaX adopts the keyboard configuration from the initial installation. Select the keyboard model and language, if necessary. You can use the test field to enter some text for testing purposes. If your keyboard is working fine, click **Next >>** to continue to the video card configuration window.

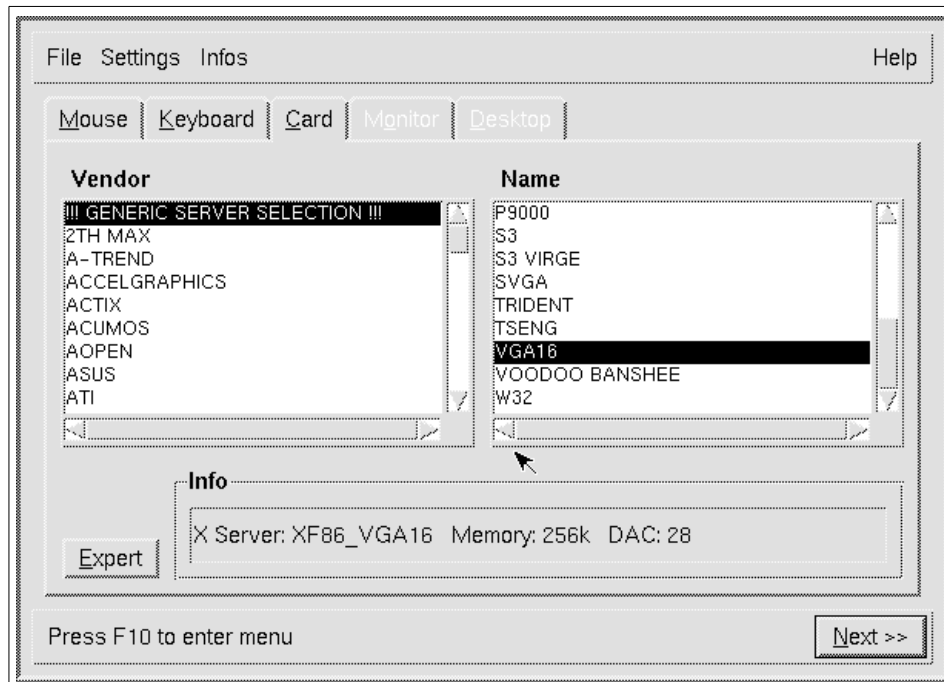


Figure 109. SaX: X Server selection

If SaX was able to detect your video card, you will see the word `Autodetect` in the Name field. You can then proceed to the monitor configuration window immediately. If your video card has not been detected, you can either select it from the Vendor list, or choose **Generic Server Selection** and select the correct X server for your video card. Some cards require additional configuration options. Click **Expert** to open the advanced configuration options dialog. Click **Next >>** to select your monitor.

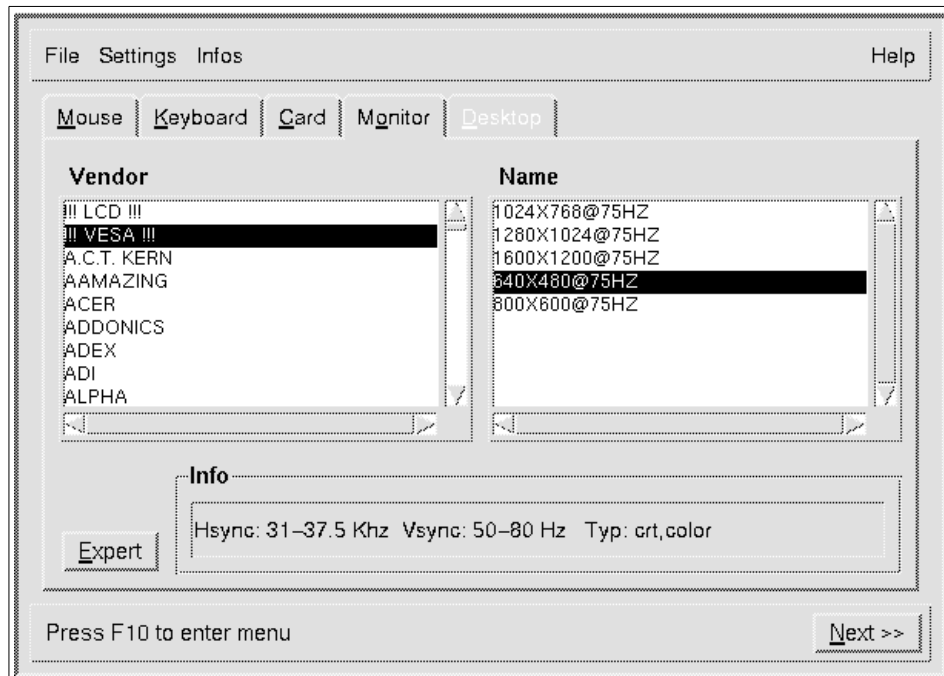


Figure 110. SaX: Monitor configuration window

In order to create an optimized screen resolution and refresh rate, SaX needs to know the capabilities (the horizontal and vertical frequency range) of your monitor.

Select your monitor vendor and name from the list. If you cannot find your model, you can either choose a generic VESA model, or enter the correct frequency range by clicking **Expert**. Please see the technical documentation of your monitor for the correct values.

Stop

Please make sure to enter the correct frequency range in the expert mode field! You can severely damage your monitor by choosing a frequency range that is too high for your model, if your monitor does not have a self-protection circuit.

Click **Next >>** to advance to the next window.

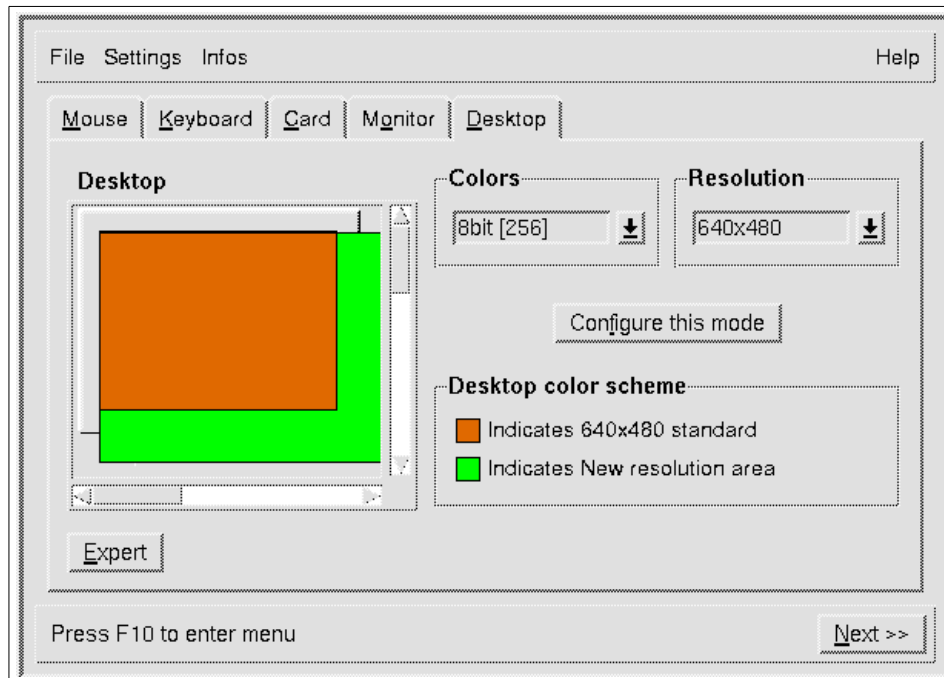


Figure 111. SaX: Screen selection window

The screen selection dialog window gives you the opportunity to select the desired color depth and resolution. SaX will only display those resolutions and color depths that will fit into your video card's memory. After choosing the correct values, click **Next >>** to test this screen resolution.

SaX will now compute the best refresh rate for this resolution and switch to the display. If you do not see a picture, your monitor powers off, or begins to flicker, press **Ctrl+Alt+Backspace**, to return to SaX. If your monitor displays the higher resolution correctly, you can now make some fine adjustments to this resolution. Click **Save** if you are satisfied with the result.

Chapter 3. Basic system administration

This chapter will give you an overview of how to perform the most common administrative tasks on a SuSE Linux system. Most of these tasks can be done with YaST, SuSE's configuration and administration tool. However, you may still edit the different configuration files manually, if you wish. YaST will detect manual changes and will not overwrite them.

3.1 Adding and removing software packages using YaST

SuSE Linux uses the RPM package manager to manage software packages of the distribution. RPM uses a database to store information about all files that belong to a certain package, including some additional information about the package. RPM itself is a command-line program. You can use it from the command line to add, remove or obtain information about software packages and system files. See 3.2, "Package management using RPM" on page 117 for more details. YaST, SuSE's administration and configuration tool, can act as a user-friendly front end to RPM.

To install or remove software packages, insert the first CD-ROM and start SuSE's installation and configuration tool YaST by typing `yast` at the command line (as user `root`). YaST will start up and you will see YaST's main menu.

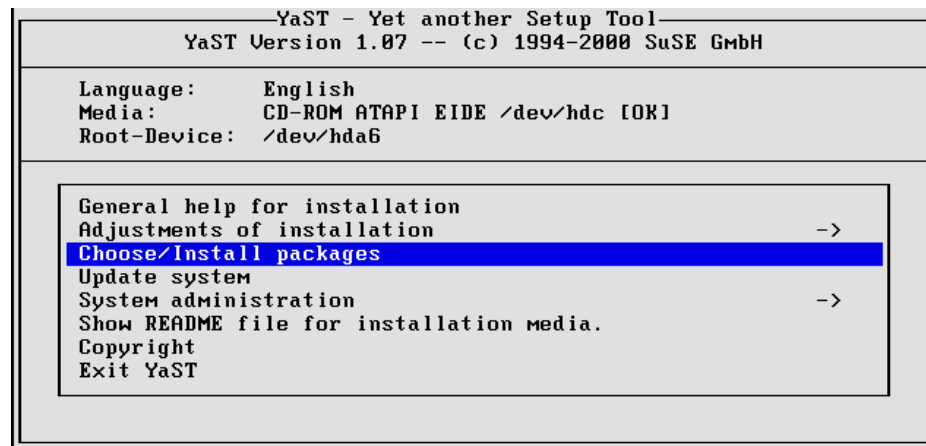


Figure 112. YaST: main menu

Highlight the menu entry **Choose/Install packages** and press Enter. Alternatively, you can invoke YaST with the following parameters:

```
yast --mask install --autoexit
```

This will automatically open the installation main menu and will return to the command line on exit.




Figure 113. YaST: package installation main menu

SuSE Linux offers a choice of software configurations. These contain a list of selected software packages to fit a certain need. Select **Load configuration** to load a predefined configuration.

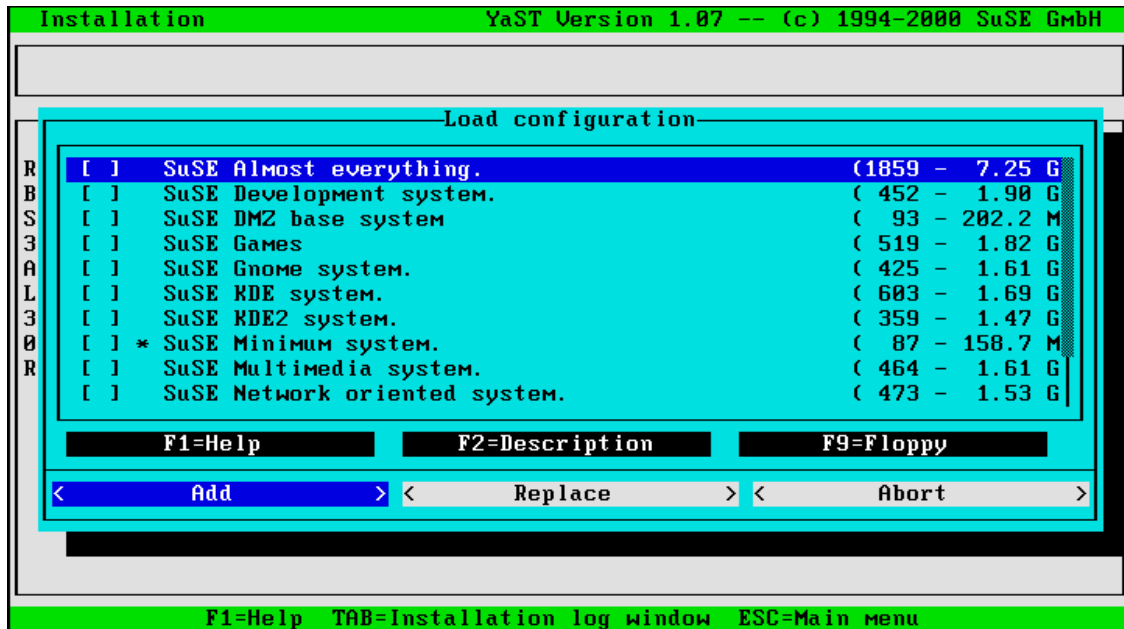


Figure 114. YaST: load software configuration

You can now add the files from a configuration to your current configuration, or you can replace it with one of these configurations. If you replace a configuration, all currently installed packages that are not part of the selected configuration will be marked for deletion! Press Esc to return to the main menu.

To add packages to or remove packages from your current configuration, select **Change / create configuration**. This will open the Series selection window shown in Figure 115.

```

Series selection                               YaST Version 1.07 -- (c) 1994-2000 SuSE GmbH
Series
doc  Documentation                             [144.7 M]
e    Emacs                                    [ 31.7 M]
emu  Emulators                                [   0 B]
fun  Games and more                           [  1.9 M]
gnm  GNOME - GNU Network Object Model Environment [   0 B]
gra  All about graphics                       [ 59.9 M]
ham  Amateur Radio (AX.25, CW, Logs, etc.)    [   0 B]
k2de KDE2 - K Desktop Environment (version 2) [   0 B]
kde  K Desktop Environment                   [ 76.8 M]
kpa  KDE applications                         [  6.5 M]
n    Network-Support (TCP/IP, UUCP, Mail, News) [ 54.2 M]
pay  Commercial Software                     [355.3 M]

<F3>=Zoom
device-name partition total   used   free   free% mount-point
/dev/hda6   Linux      1.85 G  1.12 G 753.3 M 39% /
/dev/hda5   Linux        5.5 M 325.0 K  5.2 M 94% /boot

F1=Help  F4=Resorting  F5=Dependencies  F10=Esc=Exit

```

Figure 115. YaST: series selection

All software packages are categorized into different series. Choose your category and press Enter to see all packages belonging to this series.

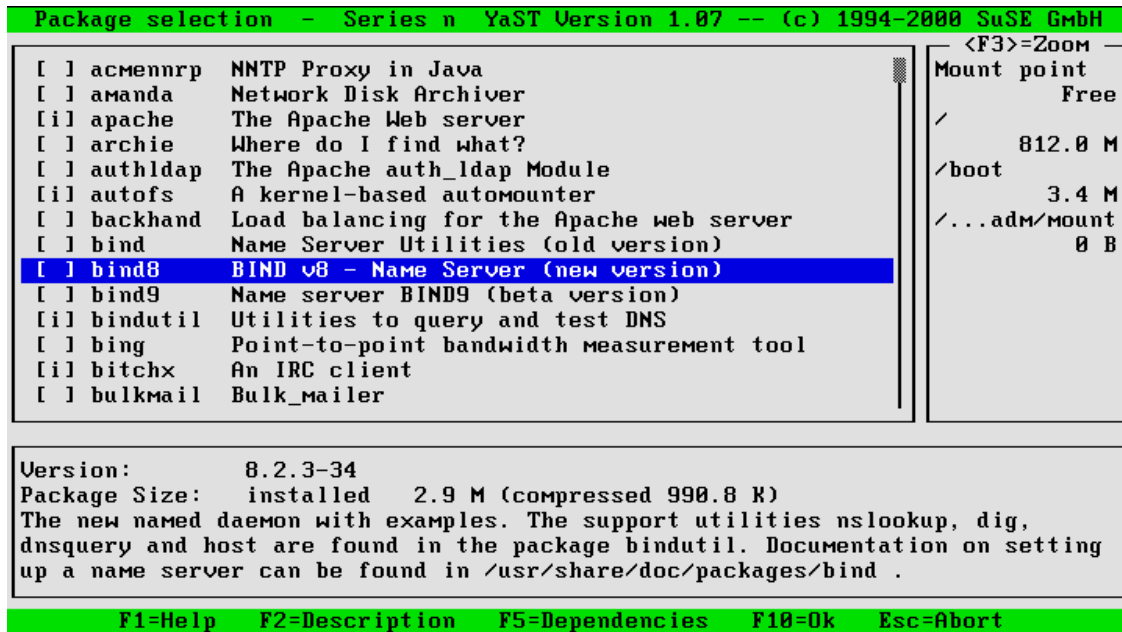


Figure 116. YaST: package selection

To select a package for installation/removal/update, press the Spacebar or Enter. This will toggle the status of the selected package. The indicator in the first column displays the current status:

Table 8. Package selection indicators

Indicator	Package status
[]	Package is not installed and not selected for installation
[X]	Package is marked for installation
[i]	Package is already installed
[R]	Package is installed and will be replaced / reinstalled
[D]	Package is installed and marked for deletion

If you want to change the package status of multiple packages at once, press Shift+A (see Figure 117).

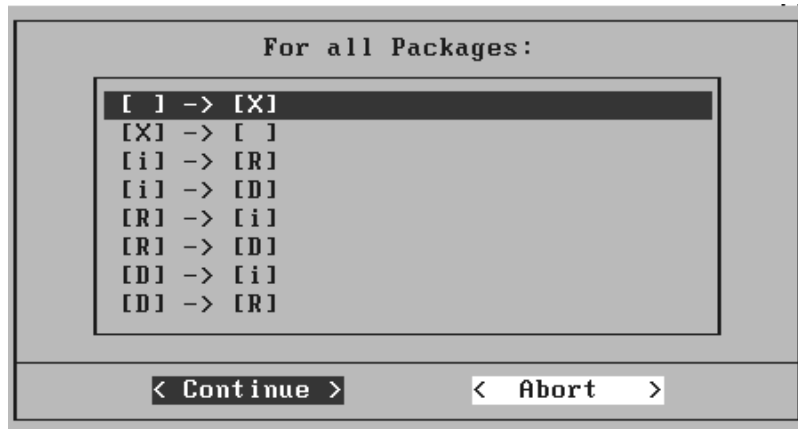


Figure 117. YaST: apply changes to all packages

After you have made your choice, press F10 to return to the series selection. You can now select or remove packages from other series, or press F10 once more to return to the software configuration main menu. If you made any modifications to your current software configuration, you can start the actual installation or removal of packages by selecting **Start Installation**. If you want to verify what packages will be installed, removed or replaced, select **What if...**

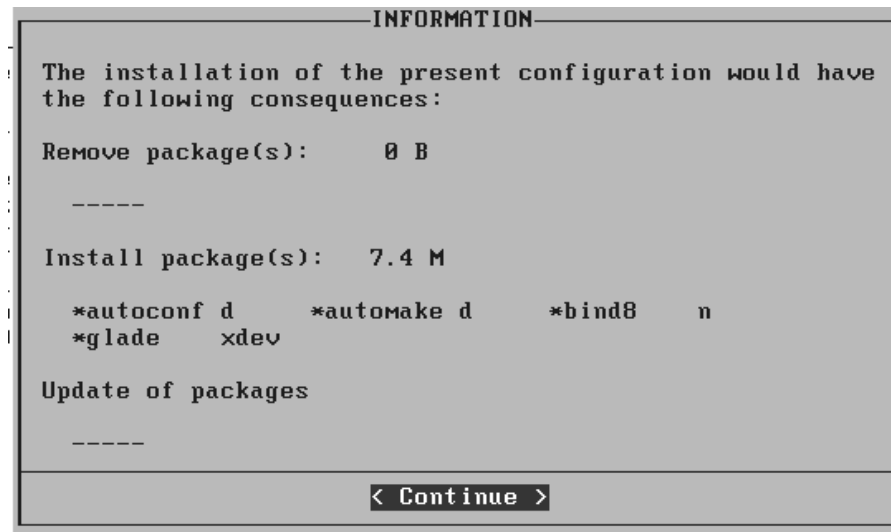


Figure 118. YaST: what if...

Click **Continue** to return to the main menu. If you are satisfied with your selection, select **Start installation**. YaST will now check on which CD the necessary packages are located and will prompt you for the respective CD. After the packages have been installed, you will return to the main menu shown in Figure 113 on page 112. You can now either add or remove additional packages. If you want to save your current package selection (for example for copying it to another system), select **Save configuration**. You will then be prompted where you want to save the configuration to. Select **to floppy** or **to hard disk**, depending on your needs. If you are saving to a floppy disk, make sure that it does not contain valuable data! The diskette will be erased during this process.

You can return to the YaST main menu by selecting **Main menu**.

3.2 Package management using RPM

Package management can also be done directly with the Red Hat package manager (RPM) on the command line. The following table shows some of the most frequently used commands.

Table 9. Basic RPM commands

Command	Description
<code>rpm -q <package></code>	If package is installed, check version and build number of installed package
<code>rpm -qi <package></code>	Obtain some more information about an installed package
<code>rpm -qa</code>	List all installed packages
<code>rpm -qf <filename></code>	Determine the (installed) package that <filename> belongs to
<code>rpm -Uhv <package.rpm></code>	Update/Install the file package.rpm showing a progress bar
<code>rpm -F -v ./*.rpm</code>	Update (freshen) all currently installed packages using the RPM files in the current directory
<code>rpm --help</code>	Get some help about the different options and parameters

Note

If you install packages using RPM on the command line, make sure to run the script `SuSEconfig` afterwards! Some packages require post-installation maintenance.

More information and options about RPM can be found in the manual page (`man rpm`), the RPM how-to

```
(less /usr/share/doc/howto/en/RPM-HOWTO.txt.gz)
```

and on the RPM Web site at <http://www.rpm.org>. You can also display a short overview by running

```
rpm --help.
```

3.3 User and group administration using YaST

Linux is a multi-user operating system. To differentiate between the various users, each user has to log in with a unique user name and password. Each user belongs to a primary user group, but they can also be a member of other groups as well (up to 16 groups). Each user name is associated with a user ID (UID), which is also unique throughout the system. The same applies to user group names and group IDs (GIDs).

Usually each user has a personal home directory. This is space on the file system (usually a directory below `/home`, for example `/home/username`) that belongs to a person and where the person can store their personal files (for example e-mail or text documents). Other users generally have no access to the files stored in another user's home directory.

You should carefully consider adding user groups before adding users. Sometimes there are concerns about restricting access to some parts of the user file system. You can do this by creating separate user groups to control access to various files and filesystems. Also if you are going to be creating a system with many users, you should consider creating separate groups divided by what they are doing on the system. You can create an admin group for admins, a `db2user` group for DB2 users, and so forth. Linux allows you to control access to both files and directories by users, groups, and everyone on the system.

Another concern in setting up users and groups is that you may want to share files with other systems. This can be done by CD-ROM, tape, diskette or any similar device. You can use the network to share information with NFS,

Samba, IPX and other network packages. If you use user and group names and characteristics that are not the same on all systems doing the sharing, then you can have file sharing and access problems.

If you are creating logins and groups on each box separately, it is often best to use a single system where all your IDs can be created. This system is then used as a reference. It is not necessary that everyone actually log into the reference system. It only exists to coordinate ID and group creation and to prevent non-standard IDs and groups. A user also cannot log into the reference system if the password is not enabled. This will prevent unauthorized access to the system. If you want to administer a lot of users on different machines, you should consider setting up NIS. See Chapter 12, “NIS - Network Information System” on page 237 or Chapter 13, “LDAP - Lightweight Directory Access Protocol” on page 247 for more information about this.

It is one of the root user’s tasks to add and remove user or group accounts. With YaST, SuSE provides an easy-to-use tool for user and group administration. To use it, log in as the root user and type the command:

```
yast --mask user --autoexit
```

Alternatively you can invoke YaST by simply typing `yast` and choosing **System administration -> User administration**. The following window will appear:




Figure 119. YaST: user administration main window

To add a new user, fill in the blanks. The user name should be short and in lowercase (YaST will do some verification on the input). After you pressed Tab or Enter to advance to the next input field, YaST will automatically look for the next available user ID and will assign it to this user. The entries Group, Home directory and Login shell will also be filled with default values, but you are free to change them to fit your requirements.

Some information about the different shells:

- **/bin/bash** - This is the Bourne Again Shell, which is an extension to the Bourne Shell. This is the most popular shell for Linux.
- **/bin/sh** - This is the standard Bourne Shell that has been around since almost the beginning of UNIX.
- **/bin/ash** - This is another version of the Bourne Shell.
- **/bin/bsh** - This is the same as /bin/ash to which it is linked.
- **/bin/ksh** - This is the standard Korn shell that is the most popular shell for UNIX Administration.
- **/bin/tcsh** - This is a public domain extension of the C Shell.

- **/bin/csh** - This is the standard C Shell that originated at the University of California, Berkeley.
- **/bin/zsh** - This is another extension of the Bourne Shell.

Your choice of shells is a matter of preference, but generally UNIX admins prefer Bourne or Korn Shell programs, whereas programmers tend to prefer C Shell-based programs.

If you want this user to be able to connect to the Internet using a modem, check **Access to modem permitted**. This will add this user to the user groups dialout and uucp, which have the necessary permissions to initiate a dial-up connection using the tool wvdial. The entry fields User name, Group and Login shell also provide a selection list where you can choose a previously defined value. Press F3 in the respective entry field.

After you have filled in all fields, press F4 to actually create the user. If the home directory of that user did not exist before, it will now be created and the contents of the directory /etc/skel will be copied into it. This skeleton directory contains a basic framework of configuration files for the user to start from.

If you want to remove a user account, just select the login name using F3 or enter the name manually in the user name input form. To delete this user, press F5 and confirm the following question with **Yes**. You will be prompted for a confirmation before the user's home directory will be removed, too.

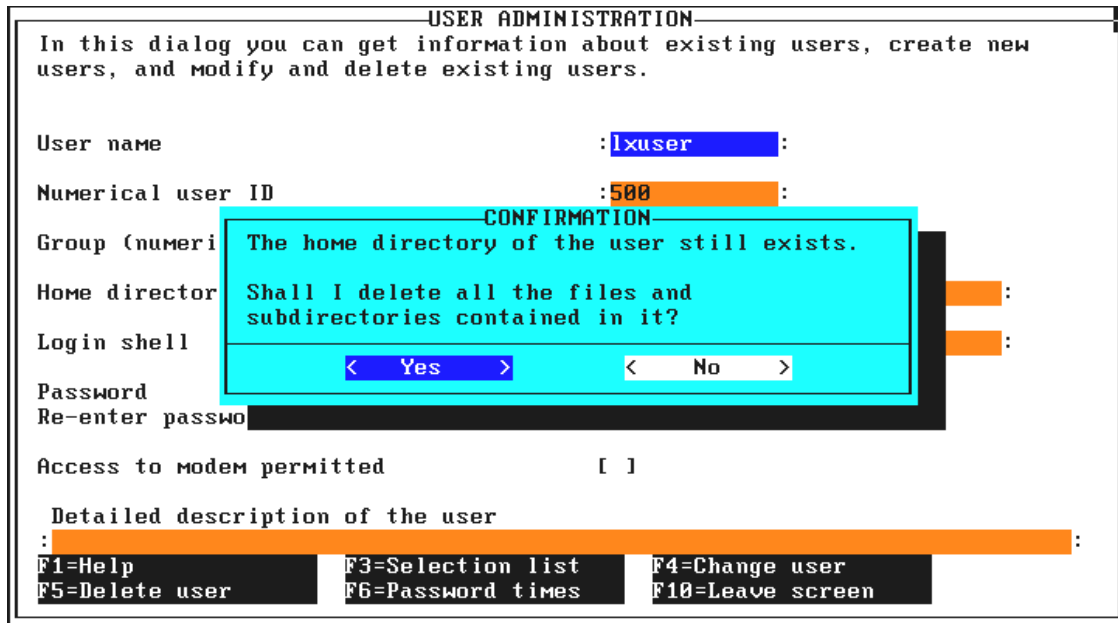


Figure 120. YaST: home directory removal confirmation

After you have finished the user administration, press F10 to return to the main menu.

3.4 Adding users on the command line

To add users to the Linux system you can also use the command `useradd`. In Linux you can find the options to `useradd` by typing the command by itself as in Figure 121. This is recommended only for commands that you know require an option. Otherwise, you may inadvertently execute a command you do not want to.

```

SuSE:~ # useradd
usage: useradd [-u uid [-o]] [-g group] [-G group, ...]
             [-d home] [-s shell] [-c comment] [-m [-k template]]
             [-f inactive] [-e expire ] [-p passwd] name
useradd -D [-g group] [-b base] [-s shell]
             [-f inactive] [-e expire ]

```

Figure 121. The `useradd` command

You can also use the `man` command to obtain more detailed information about the different parameters.

Other commands have information presented by using the `--help` option. This option is not implemented in all commands but in the case of the `useradd` command it will present basically the same information you see in Figure 121.

You can find out what your current default values are with the command `useradd -D` as shown in Figure 122.

```
SuSE:~ # useradd -D
GROUP=100
HOME=/home
INACTIVE=0
EXPIRE=10000
SHELL=/bin/bash
SKEL=/etc/skel
```

Figure 122. Default values for creating a user ID

The explanation of the options are as follows:

`-c comment`

This is a comment field about the user. It has been traditionally called the General Electric Comprehensive Operating System (GECOS) field and can include such information as office room numbers, phone numbers, etc. Any string of characters must be put into double quotes. For example, `-c comment "John Doe, rm. 45, x 78965"`.

`-d home_dir`

The home directory location of the user. If this is not specified then the default is to append the login name to the end of the default value for HOME shown in Figure 122. For example, the home directory for `jdoe` will be `/home/jdoe` unless specified here.

`-e expire_date`

This is the date on which the user account will be disabled. The date is specified in the format `MM/DD/YY` where `MM` is the month, `DD` is the date and `YY` is the two-digit format of the year. (Note that even though the date is represented in two digits, Linux converts the date to a format that is not Y2K dependent, so there are no Y2K worries here.) The default is the value of EXPIRE in Figure 122.

`-f inactive_time`

This gives the status of the account. The value of 0 says to disable the account when the password expires. A value of -1 says not to disable it. The default is the value of INACTIVE in Figure 122.

`-g initial_group`

The initial group that a user logs in with. This can be a name or number of a currently existing group. This is specified in the `/etc/passwd` file as the GID or Group ID value. The default group is given by the value of GROUP in Figure 122.

`-G group[, ...]`

This is a list of any additional existing groups the user may belong to. Each group is separated by a comma.

`-m [-k skeleton_dir]`

The `-m` option says to create the user's home directory if it does not exist. The `skeleton_dir` is the location of files that are copied to a new user's directory. The default location, if you do not use the `-m` option, is the `/etc/skel` directory. The default is the value of SKEL in Figure 122.

`-s shell`

This is the shell that the user will first log in with. The default is the value of SHELL in Figure 122.

`-u uid [-o]`

This is the numeric UID or user ID number that is used by Linux to distinguish one user from the other. All UIDs must be unique unless the `-o` option is used. The `-o` option is often used for creating IDs that have the same access rights, but different logins and passwords. The system looks only at the UID and GID values for determining access rights.

`-r`

This is used to create a system account whose UID is lower than a certain number defined in `/etc/login.defs`. You will also need to specify the `-m` option if you want to create the home directory. Otherwise, it will not be created. System accounts generally have UID values between 0 and 99.

`login`

This is the login name that the user will log in with. This will need to be unique on the system.

3.4.1 Modifying users - the command line version

You can modify user logins with the `usermod` command.

```
# usermod
usage: usermod [-u uid [-o]] [-g group] [-G group,...]
              [-d home [-m]] [-s shell] [-c comment] [-l new_name]
              [-f inactive] [-e expire ] [-p passwd] name
```

Figure 123. The `usermod` command

The options for the `usermod` command are basically the same as those for the `useradd` command, so they will not be repeated except for those that are different. With the `usermod` command you need to observe the following options.

`-d home [-m]`

The `-m` option says to move the contents of the current home directory to the new home directory and create the directory if it does not exist.

`-l new_name`

This allows you to change the user's user name that he logs in with. The user cannot be logged in with this name when he does this.

`-p passwd`

This allows you to set the password of the user from the command line. This can be useful if you have a program that automates password creation, since you can use a variable in the place of the `passwd` string.

3.4.2 Deleting users - the command line version

The command to delete users is `userdel`. You can see the options in Figure 124. This command is a lot simpler because there is not much choice you have when deleting a user.

```
# userdel
usage: userdel [-r] name
```

Figure 124. The `userdel` command

The only option that you can use is:

`-r`

This says for you to remove the home directory and its contents. Otherwise, the home directory and its contents will not be deleted.

3.4.3 Group administration using YaST

To administer user groups, select **System Administration -> Group administration** from the YaST main menu. Alternatively, start YaST from the command line using the following parameters:

```
yast --mask group --autoexit
```

This will get you directly to the group administration window:

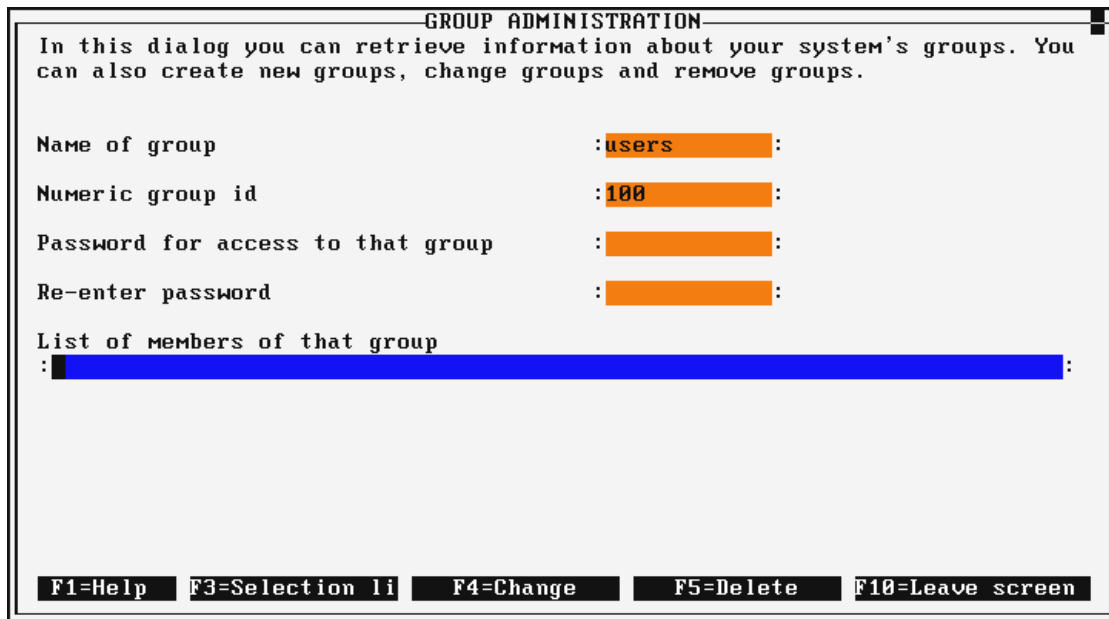


Figure 125. YaST: group administration window

Each user group has a unique name and ID. The default group for normal users is users. To create a new group, enter the name of the group and press Tab to advance to the next entry field. If you entered a new group name, YaST will automatically assign the next available group ID to this group. You can accept it or modify it to your needs. If this group is not intended to be a primary (default) user group, you can protect it with a password as well. All users that should be members of this group can be entered in the line **List of members of that group** (comma-separated). You can press F3 here to select them from the user list, or you can add them manually. Press F4 to create this group, F10 or Esc to leave this window.

If you want to delete a user group, select the group name with F3 or enter it manually and press F5 to delete it. Please note that this will not delete the user accounts belonging to this group! It will only remove the group

information from the file /etc/groups. To leave the group administration window, press F10 or Esc.

3.5 Network configuration with YaST

A Linux system will in most cases be connected to one or more networks. YaST also offers configuration options to set up your network connection. If you need to connect your host to an Ethernet or token-ring network, you can use YaST to enter the correct networking parameters. If you did not define your network card during the initial installation, or if you added a new network card to your system, you first have to define the correct driver for this device. From the YaST main menu select **System administration -> Integrate hardware into system -> Configure networking device**. From the command line, type the following command to open the network device selection window shown in Figure 127 directly.

```
yast --mask netcard --autoexit
```

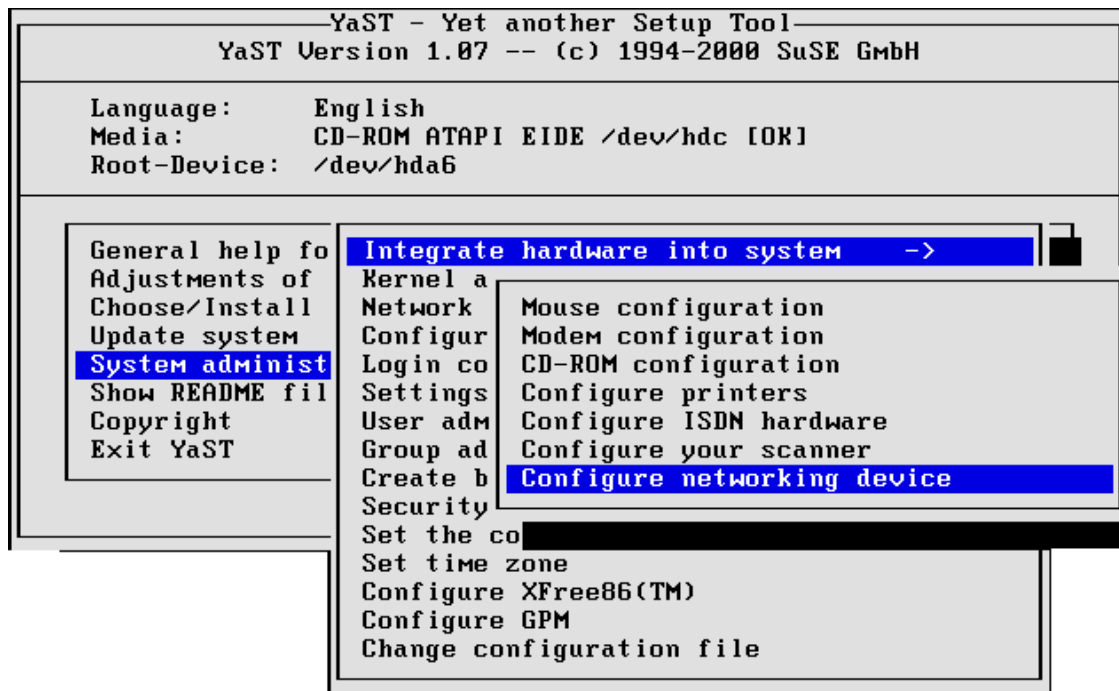


Figure 126. YaST: integrate hardware into system

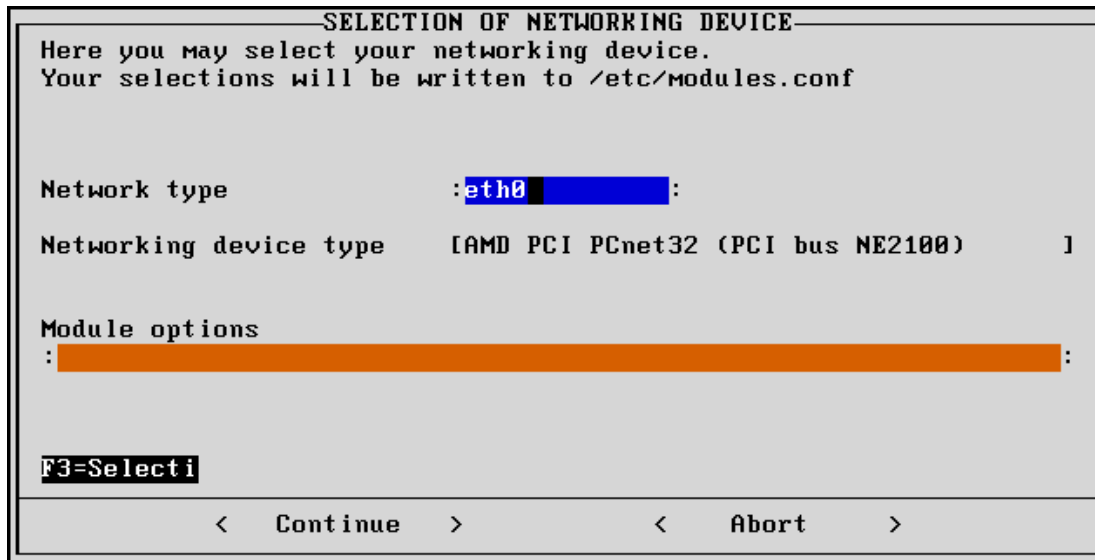


Figure 127. YaST: network device selection

First enter the network type. The two most common ones are Ethernet (for example eth0, eth1, etc.) and token-ring (for example tr0, tr1, etc.). After entering the network type, select the correct driver for this card. Some drivers need additional options; please see Chapter 14, “Kernel parameters” in the SuSE manual for a detailed explanation of the possible values. Most modern PCI network cards do not need any additional parameters, so you can most likely skip this input field. Click **Continue** to finish this configuration dialog. YaST will now add this line to the kernel module configuration file `/etc/modules.conf`.

After you defined your network type, return to the YaST System administration menu.

Now you can define the networking parameters for this device. Select **System Administration -> Network configuration -> Network base configuration**. Alternatively, type the following command at the shell prompt to jump directly to the window shown in Figure 129:

```
yast --mask network --autoexit
```

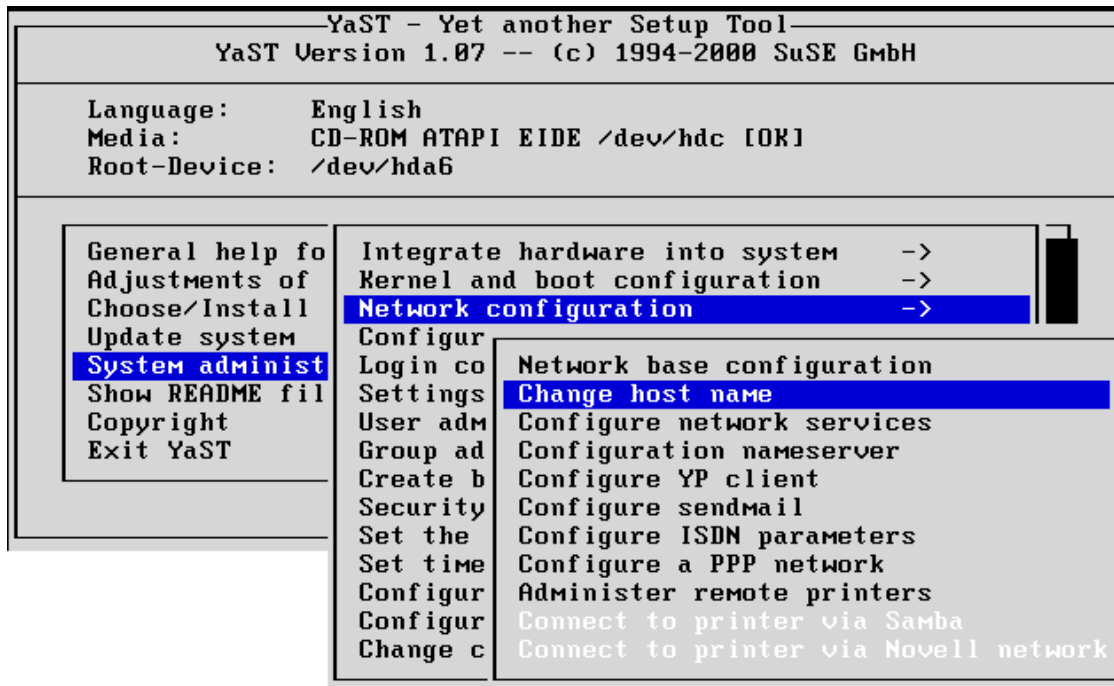


Figure 128. YaST: network configuration options

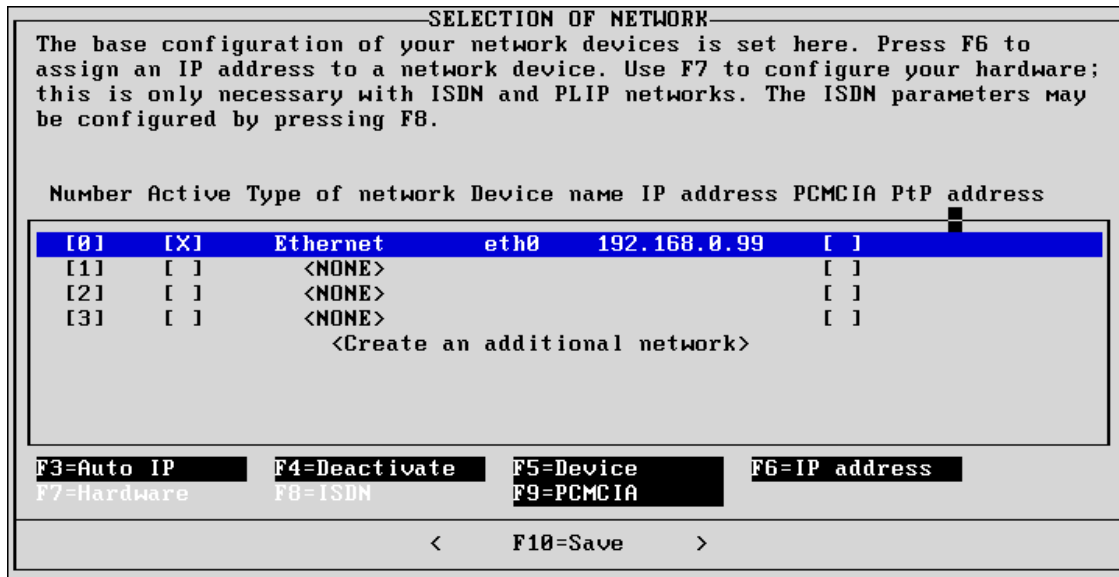


Figure 129. YaST: network base configuration

This configuration window allows you to assign IP addresses to network devices. If you have not configured your network device before, select the type of network first.




Figure 130. YaST: Set Type of Network window

Figure 130 shows the Set Type of Network selection box. Select the corresponding type for your network card and confirm the selection with **Continue**.

```
-----ENTER THE NETWORK ADDRESSES-----
Please enter the data required for the configuration of your
network. These are the IP address you want to give the machine
currently being installed (e.g. 192.168.17.42) and the netmask of
your network. The latter is 255.255.255.0 for most of the (smaller)
networks, but you may wish to set it to a different value. If you
need a gateway to access the NFS server, please enter the IP
address of the gateway host.

                                Type of network:   eth0

      IP address of your machine:  :192.168.0.99 :
Dynamic IP address (ipp, for example) [ ]
      Netmask (usually 255.255.255.0): :255.255.255.0 :
      Default gateway address (if required): :192.168.0.8 :
      IP address of the Point-to-Point partner :           :

      < Continue >                < Abort >
```

Figure 131. YaST: IP address configuration

After you have defined the network type, you can assign an IP address to this device. Press F5 to open up the dialog shown in Figure 131. Enter the IP address, Netmask and default gateway address, if necessary. Close the dialog box with **Continue**. If you configured a PLIP or ISDN device, you may also have to configure some additional hardware parameters by pressing F7.

If you have more than one network card, you can add it to the free lines below. If you need to add more than the predefined four lines, highlight **Create an additional network** and press Enter.

You can also use this dialog, if you want to assign more than one IP address to a single network card (IP aliasing). To do this, press F5 to select the type of network and choose **Enter other device**.

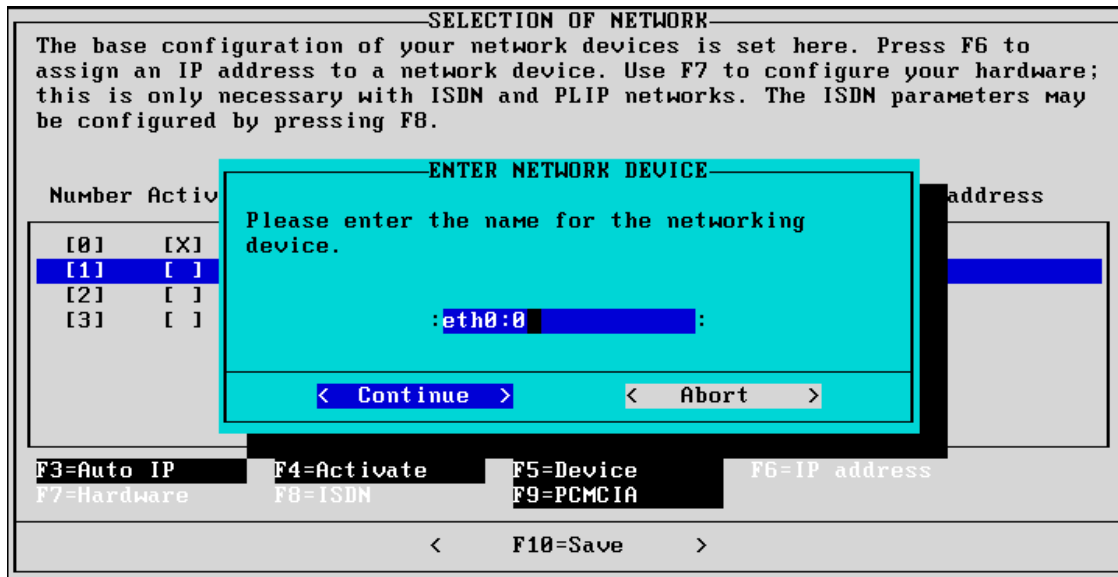


Figure 132. YaST: enter another network device

You can add multiple IP addresses to one Ethernet card, by configuring it as eth0:0, eth0:1 and so on (IP aliasing support must be activated in the Linux kernel; the default SuSE kernel has been compiled with IP aliasing support).

After you have finished the network configuration, press F10 to save the current setup. YaST will now create the respective entries in /etc/rc.config and the network setup will be applied after the next reboot or after restarting the network and routing scripts.

3.6 Changing the configuration file with YaST

SuSE Linux utilizes a central configuration file /etc/rc.config to store most of the system configuration information. The contents of this file will be used by the init scripts on bootup, as well as for creating configuration files for the different services.

The format of this file is plain ASCII text. The configuration is stored in variables in the form VARIABLE=value. Additional comments are marked with a “#” at the beginning of the line. Since rc.config contains most of the configuration information, you do not need to edit the original configuration files for most services. It is sufficient to make the change in this single file; YaST (in combination with the SuSEconfig script collection) will take care of the correct creation of these files. However, if you are used to modifying the

separate configuration files directly, you may still do so. SuSEconfig will detect the manual change and will not overwrite them. Instead you will receive a notification that SuSEconfig has detected a manual change and will create its version of this file in <filename>.suseconfig. You are free to manually implement the changes from SuSEconfig to your file.

If you want to edit variables in rc.config, you can open it in a normal text editor. Each variable has some lines of comments above its definition to give you an overview of the meaning of it. These variables are also covered in section 17.6 “The variables in /etc/rc.config” in the SuSE manual. After you have modified entries in rc.config, you have to run the script SuSEconfig afterwards to apply the changes to the different configuration files.

Alternatively, you can use YaST as a handy front end to edit these variables. From the YaST main menu, select **System administration -> Change configuration file**. To go directly to this dialog from the command line, invoke YaST with the following parameters:

```
yast --mask rcconfig --autoexit
```

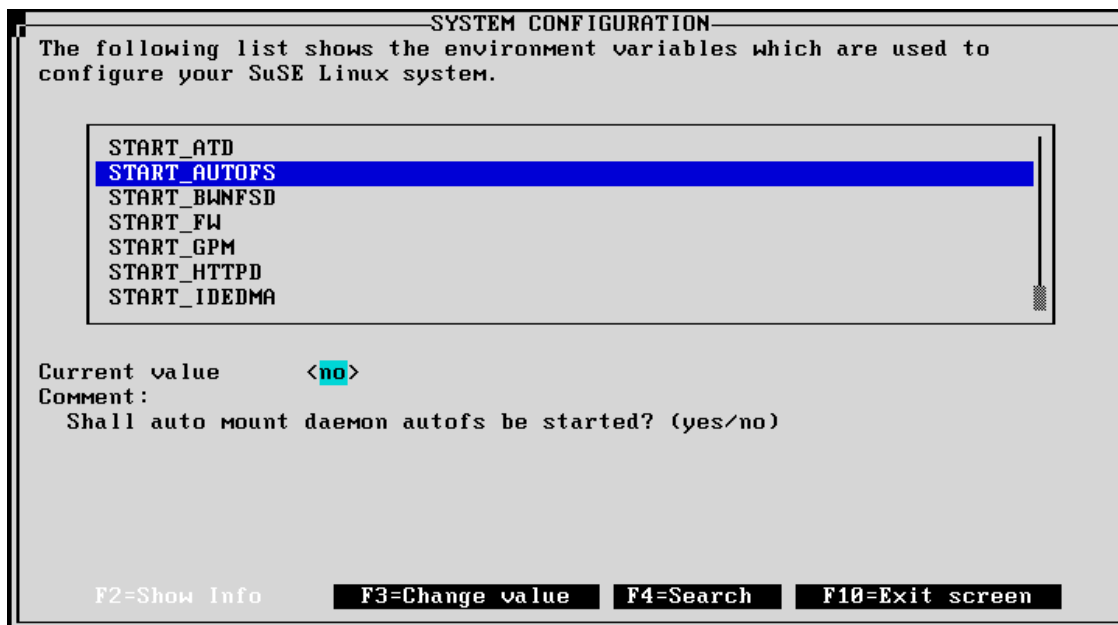


Figure 133. YaST: view the system configuration file

Use the cursor keys to highlight the desired variable. F2 gives you a description of the currently highlighted option.

To search for a certain keyword (case-sensitive), press F4 and enter the desired search term.

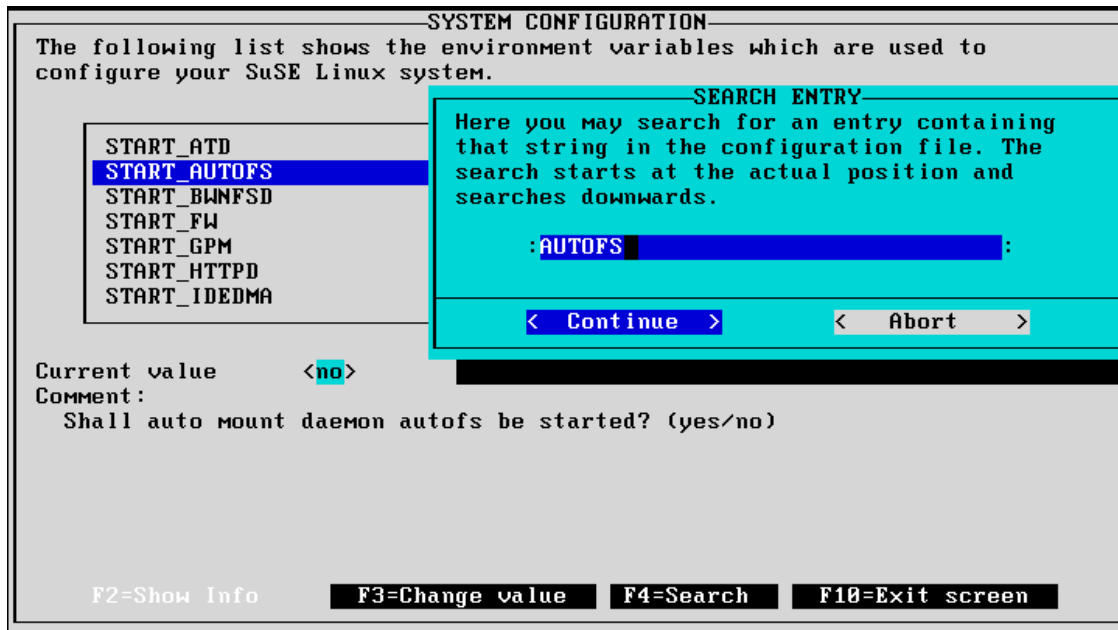


Figure 134. YaST: search for keyword in configuration file

To modify the selected entry, press F3 and enter the new value in the dialog box.

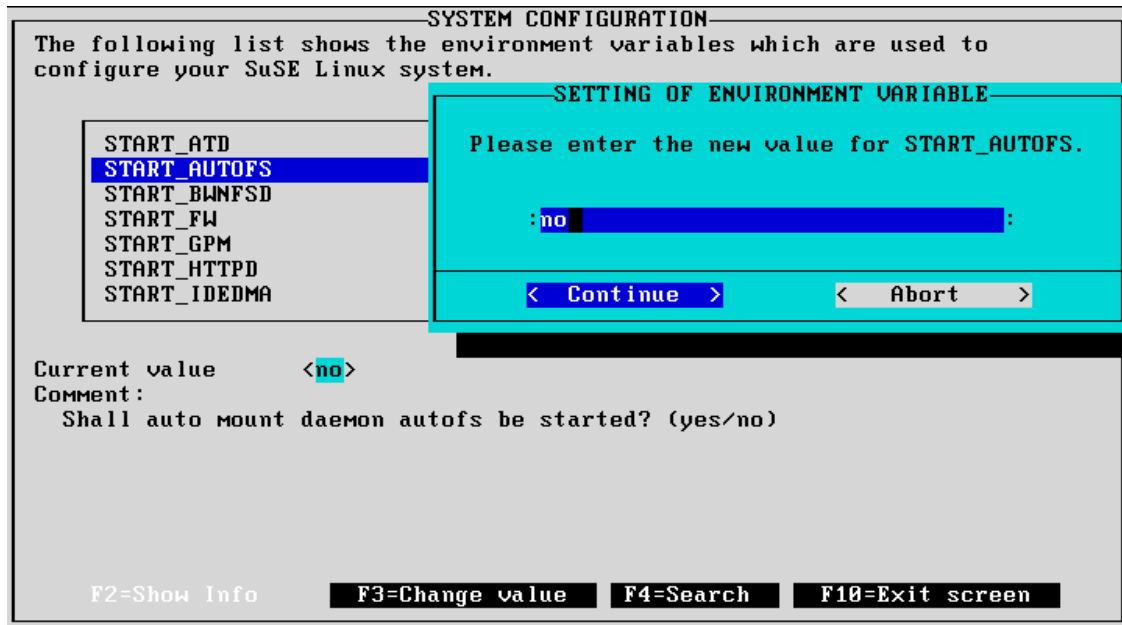


Figure 135. YaST: editing a variable in rc.config

Press F10 to finish the editing and return to the main menu after saving the changes in /etc/rc.config.

3.7 System administration with Yast2

If you prefer to use a GUI to administer your system, Yast2 is the answer. It has an easy-to-use point-and-click interface to allow first time UNIX administrators to configure a server quickly and efficiently.

Yast2 is a shell that holds a collection of modules (not to be confused with kernel modules). These modules provide the GUI component to configure a certain part of your Linux system.

Yast2 is a relatively new SuSE application, and more modules are added with every SuSE release. If Yast2 is not capable of configuring a part of the system that you wish to maintain, either use Yast1 or configure the service manually.

The SuSE technical manual details a wide range of administration procedures, and as such should be consulted if you are unsure of a procedure.

3.7.1 Yast2: Main window

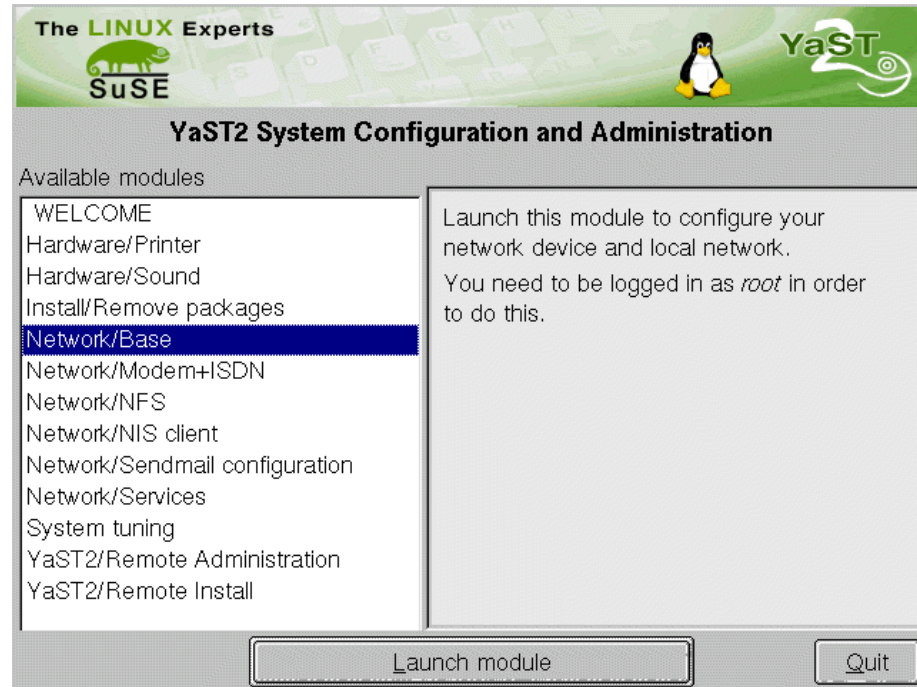


Figure 136. Yast2: Main window

Yast2 is an X-Windows application, so before you begin you need to have loaded up X before we can proceed. To load X-Windows, at the command prompt type:

```
startx
```

Note

To set up your system to boot into X-Windows instead of the console, see Section 3.6.5 Login Configuration in the SuSE manual.

Once you have loaded X-Windows, click the **Yast Menu** button on the bottom of the window (it is a gecko with a hammer and spanner behind it). Select **Yast2 - All Yast2 modules**. You will be presented with the window in Figure 136.

The left-hand pane of Yast2 shows what you can configure on your system:

- **Hardware/Printer** - Configure your printer. This can be a Novell, parallel port, USB, network, or SAMBA (Windows) printer.
- **Hardware/Sound** - Configure your sound device.
- **Install/Remove Packages** - Add and remove packages from your system.
- **Network/Base** - Configure your network settings, including network devices and IP addresses.
- **Network/Modem+ISDN** - Configure your modem or ISDN adapter, including device configuration and ISP configuration.
- **Network/NFS** - Maintain NFS exports and NFS imports for the system.
- **Network/NIS client** - Configure your machine as a NIS client.
- **Network/Sendmail configuration** - Set the behavior of sendmail on the server.
- **Network/Services** - Create, modify and remove network services from the server. This is commonly known as *inetd*.
- **System Tuning** - Speed up techniques for your system. At the moment this only allows IDE disk performance increase via UDMA settings.
- **Yast2 Remote Administration** - Allows you to administer another SuSE server from a central location.
- **Yast2 Remote Install** - Allows you to install SuSE on another computer via a serial port.

We will discuss the most useful aspects of Yast2 to a user installing on a server configuration.

3.7.2 Yast2: Network configuration

To start a module, select it from the left-hand pane shown in Figure 136 and click **Launch Module**.

Our first Yast2 module is the Network/Base configuration module. After selecting the module as detailed above you will be presented with the window shown in Figure 137.

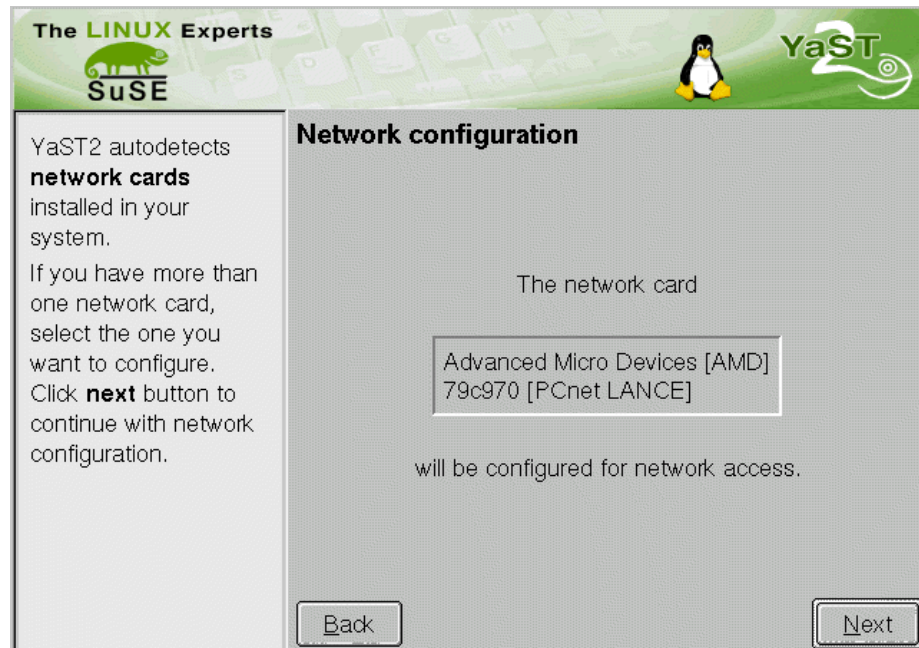


Figure 137. Yast2: Network Base configuration

Yast2 will try to autodetect your network card. If it is unsuccessful, you will have to locate the correct driver for it via the manufacturer. You may also try the SuSE FTP site to see if an update or new driver exists for the card and configure the card manually, depending on the driver.

Once your card has been detected, click **Next** to continue with the network configuration.

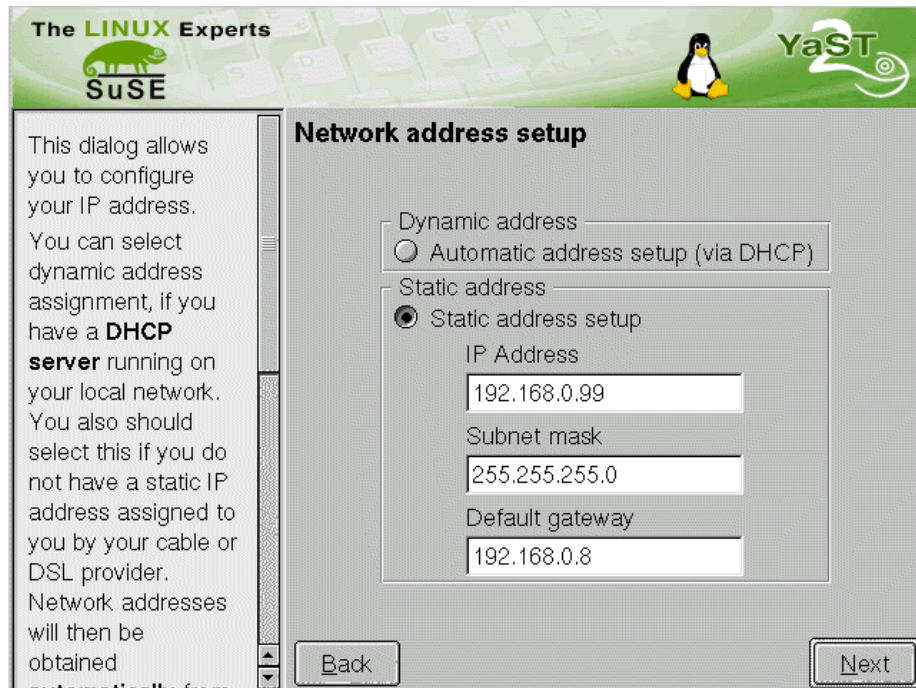


Figure 138. Yast2: Network address configuration

You have two options for configuring your network devices. You can either use DHCP to acquire your network address, gateway address, DNS server and so on, or you can manually assign an IP address to the network interface as we have done in Figure 138.

Once you have entered the corresponding values into the configuration window, click **Next** to continue.

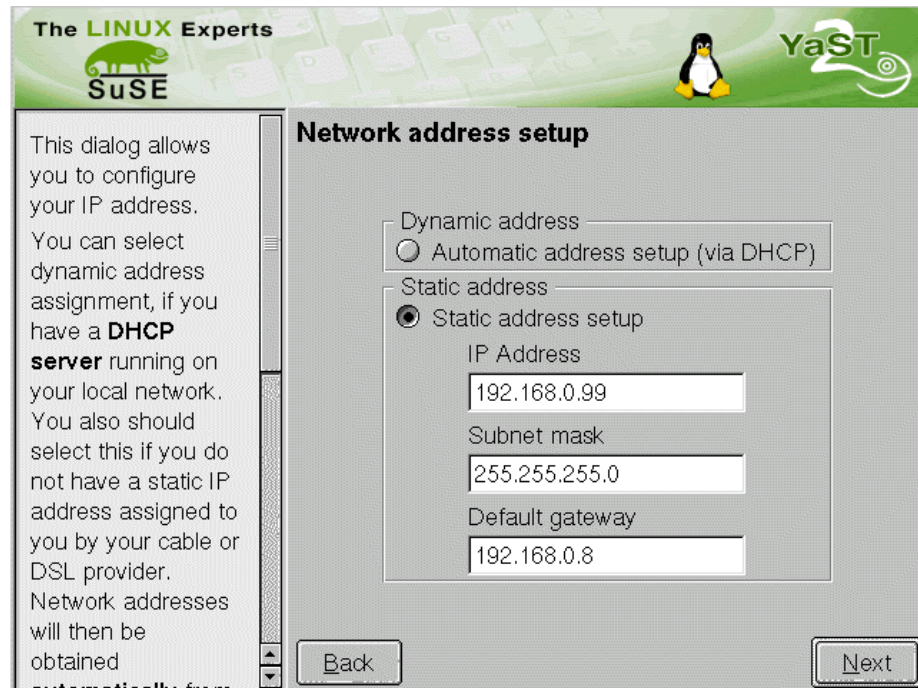


Figure 139. Yast2: Host and DNS configuration

You are now prompted to configure your host name and DNS settings. You can enter up to three domain name servers, referred to *primary*, *secondary* and *tertiary* domain name servers respectively. These should point to your DNS server, or your ISP's DNS servers. It is usual to have two domain name servers serve a domain for reasons of redundancy.

The Domain Search List refers to a wildcard list that will be appended to all non-fully qualified domain name (FQDN) names that are sent to the DNS server from this machine. For example, doing a name lookup on *netvista* will be translated to *netvista.ibm.com*, which will be queried against the name server. It allows an easy way to make nicknames for all machines on your network instead of having to type the FQDN.

Click **Finish** to complete the basic network configuration.

3.7.3 Yast2: NFS configuration

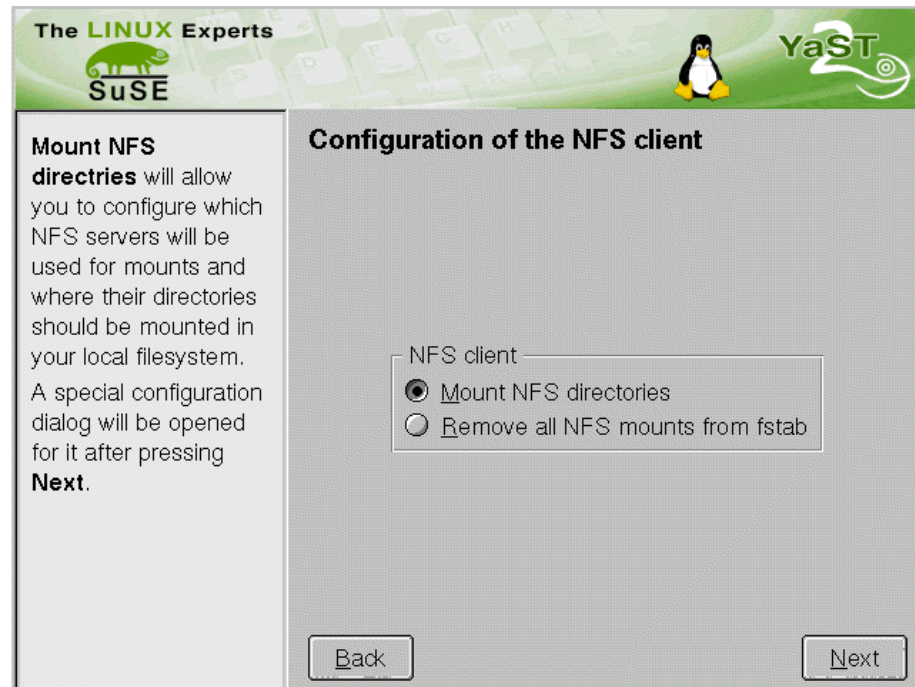


Figure 140. Yast2: NFS configuration

Yast2 allows you to configure the way NFS works on your system. This is either as a server, that allows you to share files among other machines, or as a client, that requests shares from a server. You are given the opportunity to configure both of these using the Yast2 NFS module.

To create a configuration to mount NFS shares, select **Mount NFS directories** and click **Next**. To remove all NFS mounts from your system, select **Remove all NFS mount from fstab**.

If you selected **Mount NFS directories** you will see Figure 141; otherwise you will see Figure 142.

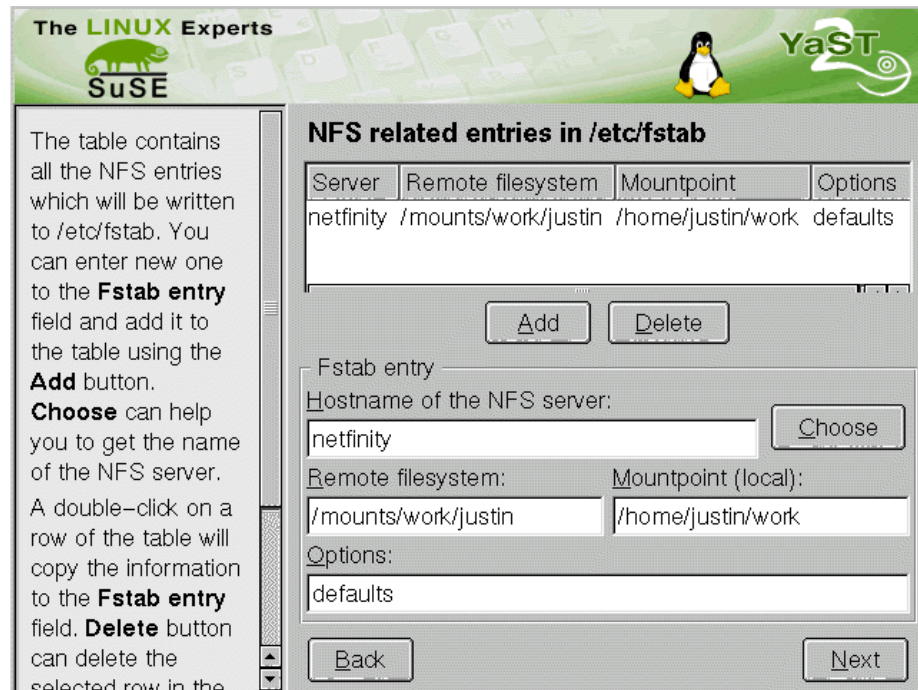


Figure 141. Yast2: Adding an NFS mount

Figure 141 is the window that allows you to configure NFS shares to be mounted by your system.

You must enter all the relevant information in this window and press **Add** to enable the share. Repeat this process until all of the NFS shares you wish to use are entered.

The entries are as follows:

- **Hostname of the NFS server** - This is the IP address or host name of the NFS server that you wish to mount the directories from.
- **Remote filesystem** - This is the remote directory on the server that you wish to request to share. It must be a fully qualified directory name, starting from the root (/) directory.
- **Mountpoint (local)** - This is the local directory that you wish to mount the remote directory under.
- **Options** - This allows you to set certain options for the mount point. Please look at the mount (8) man page for details of the options you can use.

Once you have entered all the mount points you wish to use, press the **Next** button to continue.

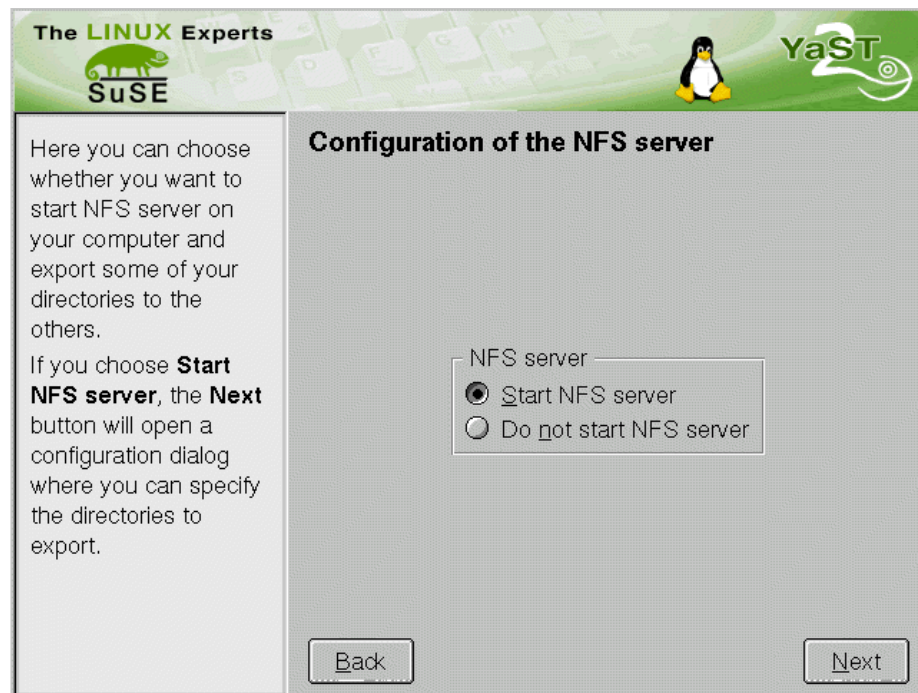


Figure 142. Yast2: Starting the NFS server

You now have the opportunity to start the NFS server to allow you to share your directories with other computers on the network. If you do not wish to share any directories with other machines, select **Do not start NFS Server** and press **Next** to continue. If you do want to share NFS mounts with other computers, select **Start NFS server**, and press **Next** to continue.

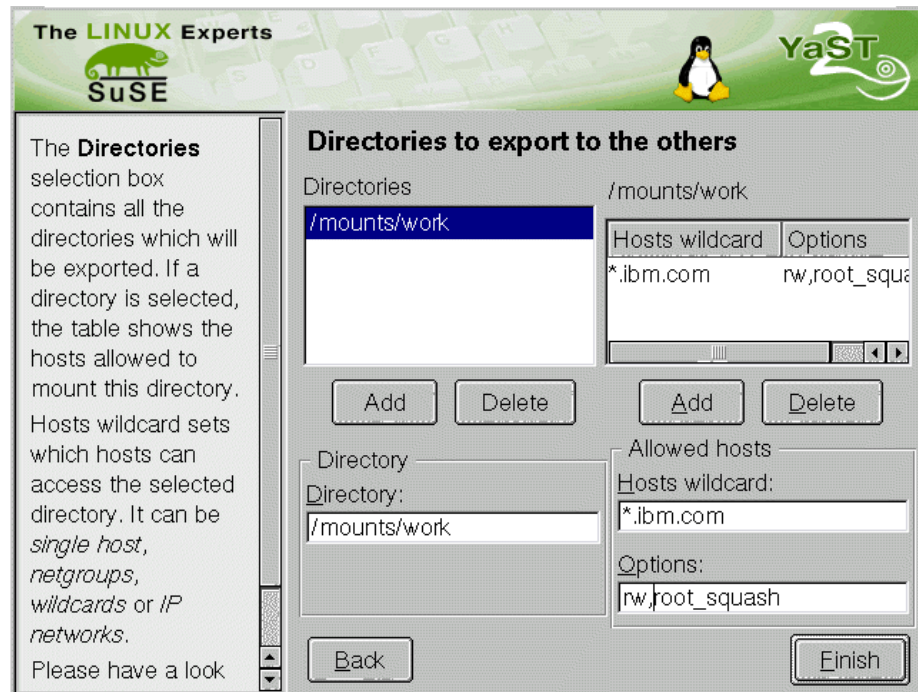


Figure 143. Yast2: Adding NFS mounts

To export your directories to other machines you have to tell the server about the directory you wish to share, who is allowed to access the directory, and under what restrictions. Figure 143 shows you how to enter this data.

As with the previous example, you must enter all the relevant data for the share, and press **Add** before proceeding with the configuration. The only difference is that you have two Add buttons.

When exporting NFS share, you are allowed to explicitly specify which hosts are allowed to access the shared directories. Yast2 allows you to keep on adding hosts that are allowed to access the specified share (in the left-hand pane). You first of all have to configure the share before imposing restrictions on who can use it. There is only one option for this, and that is **Directory**. The directory statement simply tells the NFS server what directory you wish to share. Make sure the directory exists; otherwise the server will behave erratically.

Once you have allocated a directory to share, you can start allocating its share restrictions.

The NFS server will allow host name wild cards to say a certain network can access the shares. In our case we have enabled the IBM network to access these shares by specifying `*.ibm.com` as the allowed hosts.

The options section tells the server how the restrictions impose the mount on the NFS clients. This only applies per restriction, not to every client that accesses the NFS shares. For example, if we added `*.suse.de` to the restrictions table, we could allow everyone at IBM to have write access to the share (option `rw`), but read-only access to everyone at SuSE (option: `ro`) for the same share.

See `man export(5)` for other options you can use.

Once you have added all of the mounts you wish to export, click **Next** to continue. You will be asked to confirm that you wish to use these settings. If you wish to edit them some more, click **No**; otherwise click **Yes** to commit them.

3.7.4 Yast2: Network services configuration

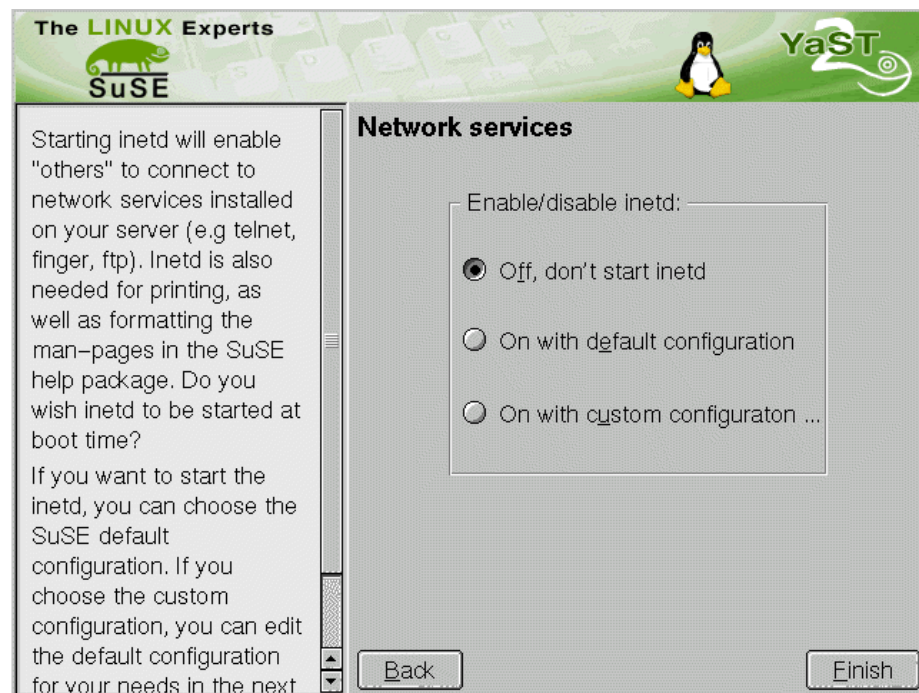


Figure 144. Yast2: Configuration of inetd

You can stop `inetd` from running at system bootup, by selecting **Off, don't start inetd**. You can use the default configuration, by selecting **On with default configuration**. Or, you can configure `inetd` yourself by selecting **On, with custom configuration...**

We will guide you through editing the `inetd` configuration to allow you to add or remove services from your server.

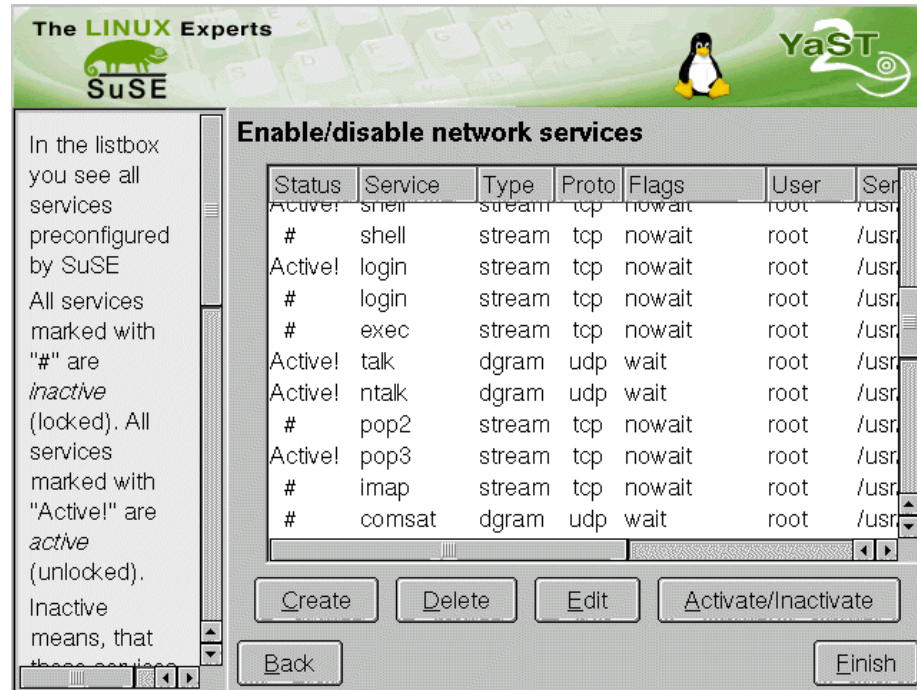


Figure 145. Yast2: Editing the `inetd` configuration

You have four options while editing the `inetd` configuration:

- **Create** - This allows you to add a new service to the server.
- **Delete** - This will delete the selected service from the system.
- **Edit** - This allows you to edit the currently selected service.
- **Activate/Deactivate** - This will stop the service, but will not delete the entry. This has the same effect as deleting the entry, but will not remove it from the configuration file. If it activates the service, it will take it out of the deactivated state and allow it to run.

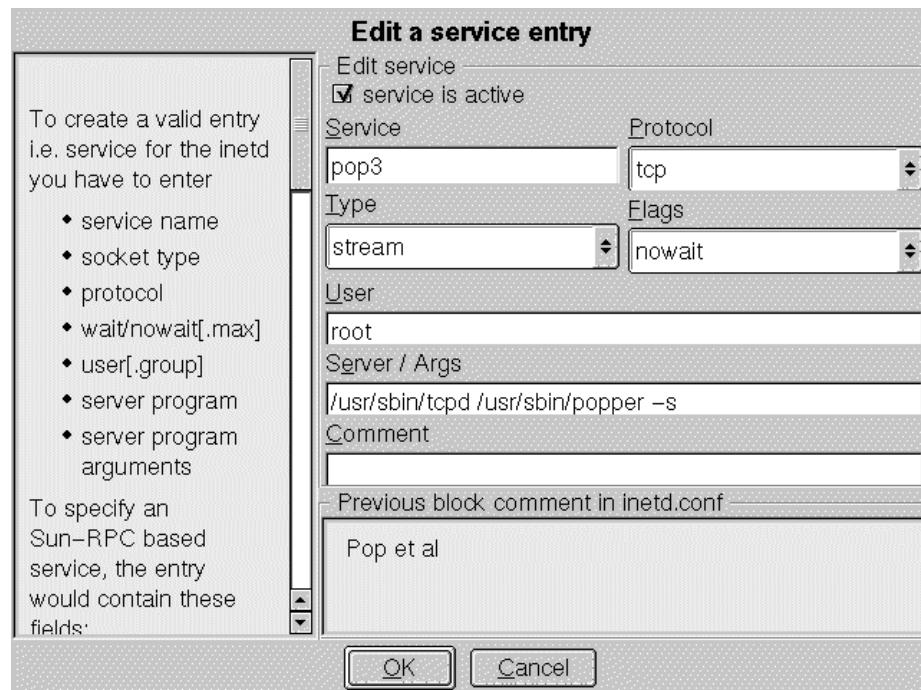


Figure 146. Yast2: Add/configure a service

Figure 146 will be loaded if you click the **Add** or **Edit** button in Figure 145. It allows you to add a service entry, or edit an existing one. The options are exactly the same for both configuration type:

- **Service** - This is a service name that is defined in `/etc/services`. This file holds information, such as port number, service type, service name and so on, regarding a certain service. You should enter an existing service name (as defined in `/etc/services`) here.
- **Protocol** - This defines what protocol this service uses. The most popular protocol types are TCP, UDP, and ICMP. The protocol defined must be present in `/etc/protocols`.
- **Type** - This defines the type of the connection that will be used. This can be one of **stream** (stream type), **dgram** (datagram type), **raw** (raw socket type), **rdm** (reliably delivered message type) or **seqpacket** (sequenced packet type).
- **Flags** - There are two options for this item. **Nowait** is usually selected for servers that use the type `stream`. It allows the service to accept new requests while processing other requests. The service is known to be “multi-threaded”. **Wait** is used to allow only one connection at a time to the

service. It is known to be “single-threaded”. Check the documentation of the service to see how it should be configured.

- **User** - This specifies under what user the service should be run. It is usually root, but it is imperative that you check the documentation of the service you are configuring, since running services under the wrong user (that user being root) can cause major security issues.
- **Server/Args** - This is the command to run the service, along with the arguments it takes. Consult the documentation of the service to find out what arguments it takes, and what those arguments do.
- **Comment** - This allows you to set a comment for this service. It is always a good idea to comment services so that you can remind yourself and others about what the service does, or special warning for other administrators.

3.7.5 Yast2: Package maintenance

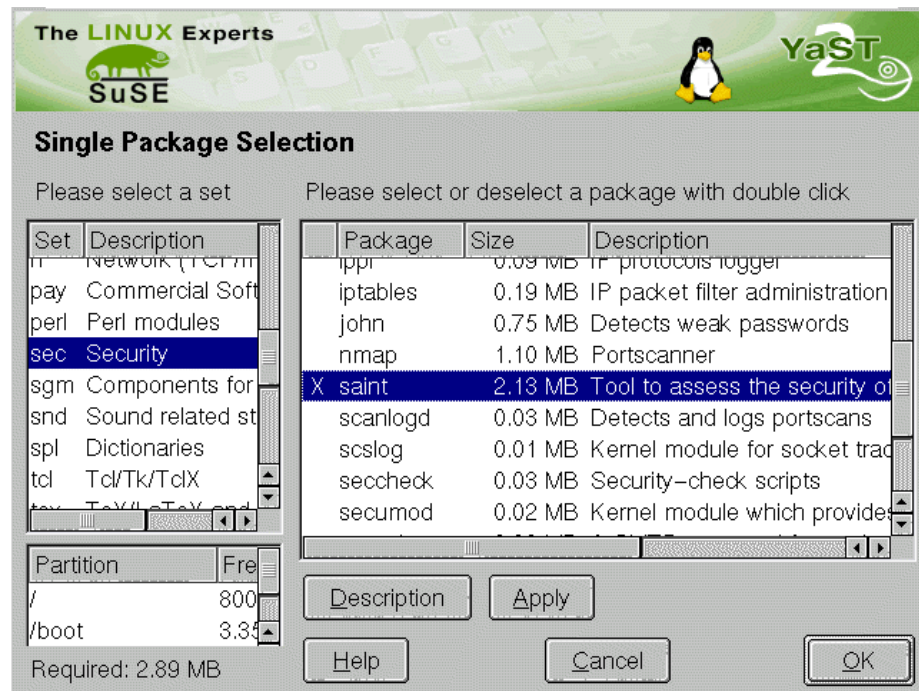


Figure 147. Yast2: Package installation

Selecting **Install/Remove packages** will allow you to install and remove packages from the system.

The left-hand pane allows you to select the package series from the installation medium. To install a package, just double-click the package name, or click **Apply** and an X will appear next to it. To remove an installed package, again double-click it: this time a d will appear next to it. This signifies that the package is marked for deletion.

To read the package description, select the package and click **Description**. This will bring up a window that will give you a short description of what the package does.

There are some combination of packages that are inadvisable to install together. If this situation arises, you will be told about the problem. The same is true if a certain package depends on other packages to run.

3.8 Finding Linux commands

You may want to run a Linux program from the command line prompt. If so, there are several directories that contain commands that you can run. You can run these without needing to know where they are because your search path includes a number of directories that will be searched whenever you try to execute a command. The search path is given by the environment variable `$PATH`. You can view the content of this variable by running the following command:

```
echo $PATH
```

If you want to find out where a command is located, execute the command:

```
whereis command_name
```

where `command_name` is the command you are looking for. If you want to find the command `yast` you can execute:

```
whereis yast
```

This will give you the following results:

```
yast: /sbin/yast
```

You notice that this command is located in the `/sbin` directory. Many of the major administrative commands will be found in the `/sbin` and `/usr/sbin` directories.

Another helpful command for finding files on your system is `locate`. The `locate` command will also list files that match the search name, if they are not

in your current search path. To search for all README documents on SuSE Linux run the following command:

```
locate README
```

Since this will be a huge amount of output, you might want to redirect the output to a text pager such as `less` or `more`:

```
locate README | less
```

This will enable you to look at the output page by page. Press `q` to leave `less` and return to the command line.

Note

SuSE Linux automatically runs `updatedb` once every 24 hours. If you cannot find what you are looking for, run `updatedb` from a command line.

3.8.1 File system permissions

Linux has inherent security features, the most noticeable being file system permissions. Setting permissions on files allows the system administrator to restrict access to parts of the file system.

File permissions can be set on files and directories. The easiest way to see an example of this is looking in the `/home` directory:

```

mail:/home # ls -l
total 1
drwxr-xr-x 19 root    root    396 Nov 15 21:06 .
drwxr-xr-x 22 root    root    467 Nov 13 16:28 ..
drwx----- 6 davej   users  912 Nov 15 21:05 davej
drwx----- 6 george  users  912 Nov 15 21:03 george
drwx----- 6 ivo    users  912 Nov 15 21:02 ivo
drwx----- 6 jakob  users  912 Nov 15 21:03 jakob
drwx----- 6 jasmin  users  912 Nov 15 21:04 jasmin
drwx----- 6 jens   users  912 Nov 15 21:04 jens
drwx----- 6 jhaskins users  912 Nov 15 21:02 jhaskins
drwx----- 6 justin  users  912 Nov 15 21:06 justin
drwx----- 6 lenz   users  912 Nov 15 21:03 lenz
drwx----- 6 linux  users  912 Nov 15 21:03 linux
drwx----- 6 malcom  users  912 Nov 15 21:04 malcom
drwx----- 6 rachael users  912 Nov 15 21:03 rachael
drwx----- 6 rafiu  users  912 Nov 15 21:04 rafiu
drwx----- 6 ruediger users  912 Nov 15 21:04 ruediger
drwx----- 6 rufus  users  912 Nov 15 21:02 rufus
drwx----- 6 ted    users  912 Nov 15 21:03 ted
drwx----- 6 uzi    users  912 Nov 15 21:04 uzi
mail:/home #

```

Figure 148. Viewing file permissions

Taking the user **linux** as an example:

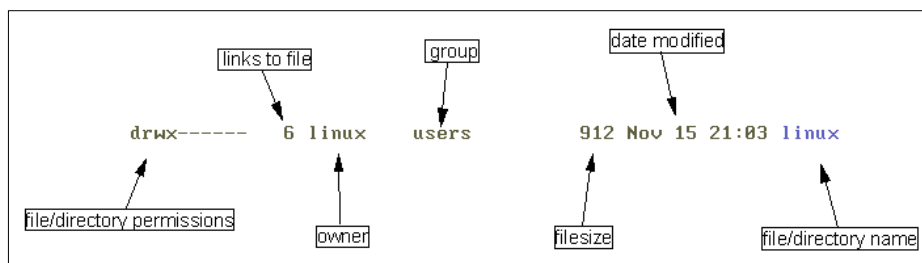


Figure 149. Explanation of ls output

What we are most interested in is the file/directory permissions. This signifies a lot of information in a short amount of space:

d - The first character in the permissions signifies that this is a directory. Other files are represented by:

- - a normal file.

l - a symbolic link to another file.

c - refers to files in the /dev directory. This signifies the file represents a character device.

b - refers to files in the /dev directory. This signifies the file represents a block device.

rwX - In this case it allows only the owner of the file (in this case linux) to read, write and execute this file.

Type	Owner	Group	World
d	rwX	---	---

As you can see, the format of the string is becoming a bit easier to understand.

The owner of the file is the user that created the file. The group part is the group that owns the file (for example, the group *users*). The world part means everyone else. Setting a permission in the world part sets the permission for every user, irrelevant of their group membership and so on.

Here is another example:

```
-rwxr-xr--
```

This means that this is a normal file, the owner can read, write and execute the file, the group can read and execute the file, and everyone else can read the file, but not modify or execute it.

As for directories, if you set a directory as:

```
drwxrwx-rw-
```

you are saying that only the directory owner is allowed to execute something “inside” the directory. So if another user tries to change directory (*cd*) into this directory, they will get a “permission denied” error message. This is exactly what happens with regards to user’s home directories.

To change the permissions on a file, you use the `chmod` command. Only *root* users can modify files that do not belong to them. You must own the file to be able to change its permissions.

The easiest way to change permissions is to use symbolic representations of what you want permissions to be.

Note

The other way to represent file permissions is to use octals. For more information about this and the `chmod` command see the `chmod` man page.

```
chmod g+rw myfile
```

This is one of the simplest ways of changing a permission. You are saying that you want the file `myfile` to allow all members of the group to be able to read and write to it.

If you used a `-` (minus sign) instead of a plus, you would be taking away those permissions. This would mean that members of the group would not be allowed to read or write to the file.

You can mix adding and removing permissions in the same command:

```
chmod u+x-rw myfile
```

This will allow executing the file, but will not allow reading or writing the file for the file owner.

Here is a summary of the symbolic representations available in `chmod`:

r - read

w - write

x - execute

- - take away the permissions

+ - add the permissions

s - set the SUID bit. This says that if the file is executable, it will be run as the owner of the file, not as the user that is running the file.

Chapter 4. Samba

If you look at any English dictionary, Samba is defined as a Brazilian dance, but Samba in Linux is something completely different. Samba is an implementation of a Server Message Block (SMB) protocol server that can be run on almost every variant of UNIX in existence. Samba is an open source project, just like Linux. The entire code is written in C so it is easily ported to all flavors of UNIX. Samba is a tool for the peaceful coexistence of UNIX and Windows on the same network on the level of file and print sharing over the NetBIOS protocol. It allows UNIX systems to move into a Windows "Network Neighborhood" without causing a mess. With Samba, UNIX servers are acting like any other Windows server, offering their resources to the SMB clients. Recently SMB was renamed by Microsoft to Common Internet File System (CIFS).

4.1 What can you do with Samba?

- With Samba, a Linux server can act as a file/print server for Windows networks. It can replace expensive Windows NT file/print server in this role, creating a less expensive solution.
- Samba can act as a NetBIOS name server (NBNS) in a Windows world, where it is referred to as WINS Server - Windows Internet Name Service.
- Samba can participate in NetBIOS browsing and master browser elections.
- Samba can provide a gateway for synchronizing UNIX and Windows NT passwords.
- With Samba client software, you can access any shared directory or printer on Windows NT servers or Samba servers and allow UNIX machines to access Windows NT files.
- With Samba File System (SMBFS) you can mount any share from a Windows NT server or Samba server in your directory structure (this is available only on Linux).

4.2 Setting up the Samba server

To see if Samba is installed, issue:

```
rpm -q samba
```

at a command prompt. If you receive the Samba version, then it is installed; otherwise refer to 3.1, “Adding and removing software packages using YaST” on page 111. Samba is in the *n* series, package name *Samba*.

4.2.1 Configuring the Samba server

In this section we will explain how to configure Samba so that it can participate as a file/print server in an existing Window network, or as a stand-alone file/print server for Windows and Linux clients.

Before you can start using Samba you need to configure the `smb.conf` file. This file is the heart of the Samba server. When the Samba package is installed in SuSE, the sample configuration file is installed as the `/etc/smb.conf` file.

The SAMBA configuration file `smb.conf` is divided into two main sections:

1. Global Settings - here you set up parameters that affect the connection parameters and settings that affect all shares.
2. Share Definitions - here you define shares. A share is a directory on the server that is accessible over the network and shared among users. This section has three subsections:
 - a. Homes - in this subsection you can define the user’s home directories.
 - b. Printers - in this subsection you can define the available printers.
 - c. Shares - this subsection can have multiple entries, one for each share you want to define.

In the following sections we will describe how to modify the `smb.conf` file to efficiently and simply use Samba as a file/print server. We explain only the most necessary parameters. If you need more information, see the manual entry for the `smb.conf` file or the Samba project Web site at:

<http://www.samba.org>

You can find our `smb.conf` configuration file in Appendix B, “Sample `smb.conf` SAMBA configuration file” on page 377.

4.2.1.1 Setting the NetBIOS parameters

The NetBIOS parameters are part of the Global Section. When you open your `smb.conf` file you will see something similar to this:

```
; Copyright (c) 1999 SuSE GmbH Nuernberg, Germany.  
;  
[global]
```

```

workgroup = Workgroup
guest account = nobody
os level = 2

```

The parameters are described in Table 10.

Table 10. NetBIOS parameters

Parameter	Description
netbios name	Although not specified in the SuSE smb.conf file, you can set the name of this server on the Windows network. This parameter has the same meaning as the Windows NT computer name. If not specified, as in our case it defaults to the host name.
workgroup	This parameter specifies in which Windows NT domain or workgroup the Samba server will participate. It is equivalent to Windows NT domain or workgroup name.
server string	Not specified in the default SuSE smb.conf, this is the description string of the Samba server. It has the same role as the Windows NT description field.
guest account	This is the Linux user that will be used to allow access to “guest shares” on the system. If a real user is assigned as the guest account, then anonymous users of the shares will have the same rights as this user.
os level	This assigns a certain level to this machine when a browser election is called. If you do not wish this machine to take part in browser elections, set this to 0.

4.2.1.2 Global printing settings

In the smb.conf file you will see something similar to this:

```

load printers = yes
printcap name = /etc/printcap
printing = bsd

```

The parameters are described in Table 11.

Table 11. Printing parameters

Parameter	Description
load printers	This parameter controls if Samba loads all printers in the printcap file for browsing.

Parameter	Description
printcap name	With this parameter you tell Samba the location of the printcap file. The default value is /etc/printcap.
printing	This parameter tells Samba what printing style to use on your server. SuSE, by default uses the BSD printing style.

4.2.1.3 Global security settings

In your smb.conf file you will see something similar to this:

```
security = user
; password server = 192.168.0.10
encrypt passwords = yes
```

The parameters are described in Table 12.

Table 12. Security parameters

Parameter	Description
security	This parameter has four possible values: share, user, server, domain
password server	In the case of server or domain security level this server is used for authorization. For the parameter value you can use the server NetBIOS name or an IP address.
encrypt passwords	By setting this parameter to yes, you enable Samba to use the Encrypted Password Protocol, which is used in Windows NT Service Pack 3 and in Windows 98. This is needed to communicate with those clients.

The default location of the Samba password file is /etc/smbpasswd.

The security modes are as follows:

- Share - for this security mode, clients only need to supply the password for the resource. This mode of security is the default for Windows 95 file/print server. It is not recommended for use in UNIX environments, because it violates the UNIX security scheme.
- User - the user/password validation is done on the server that is offering the resource. This mode is most widely used.
- Server - the user/password validation is done on the specified authentication server. This server can be a Windows NT server or another Samba server.

- Domain - this security level is basically the same as the server security level, with the exception that the Samba server becomes a member of a Windows NT domain. In this case the Samba server can also participate in such things as trust relationships.

Because Windows NT 4.0 Service Pack 3 or later, Windows 95 with the latest patches, and Windows 98 use the encrypted passwords for accessing NetBIOS resources, you need to enable your Samba server to use the encrypted passwords. Before you start the Samba server for the first time, you need to create a Samba encrypted passwords file. This can be done with the `mksmbpasswd` utility. The recommended way is to first create the user accounts in Linux and then create the Samba password file with the command:

```
cat /etc/passwd | sh /usr/lib/samba/scripts/mksmbpasswd.sh > /etc/smbpasswd
```

This creates the Samba password file from the Linux password file.

Note

By default the passwords for the Samba users are undefined. Before any connection is made to the Samba server, users need to create their passwords.

Now you need to specify the password for all users. If you are changing or specifying a password for a user, you can do this by executing the command:

```
/usr/bin/smbpasswd -U username
```

You will see a window similar to Figure 150.

```
[root@nf5000 ~]# /usr/bin/smbpasswd -U user
New SMB password:
Retype new SMB password:
Password changed for user user.
[root@nf5000 ~]#
```

Figure 150. Specifying the password for Samba user

Another way is to have each Samba user change the password for himself, by remotely connecting to the Samba server and executing the command:

```
/usr/bin/smbpasswd
```

The output will be similar to Figure 150. If a Samba user already has defined a password they will need to type the old password before they can change to a new password.

If you want to add another user to the Samba server user later, it can be done with the following command:

```
/usr/sbin/smbpasswd -a username password
```

This will add a new user to the Samba password file.

Note

You have to be logged on as root if you want to manage other users. If you are logged on as a user, you can only change your own password. The `smbpasswd` utility uses the location of the password file from the `smb.conf` configuration file.

4.2.1.4 Global name resolution settings

In your `smb.conf` file you will see something similar to this:

```
wins support = yes  
; wins server = w.x.y.z
```

The parameters are described in Table 13.

Table 13. Name resolution parameters

Parameter	Description
<code>wins support</code>	If this option is enabled the Samba server will also act as a WINS server.
<code>wins server</code>	With this parameter you tell Samba which WINS server to use.

Note

Samba can act as a WINS server or a WINS client, but not both. So only one of the parameters (`wins support` or `wins server`) can be set at the same time. If you specify the IP address of WINS server, then `wins support` must be set to `no`.

4.2.1.5 Creating shares

In the previous section we explained how to prepare general configuration parameters. But a Samba server is only of use if you offer resources to users.

In this section we will explain how to create a share. The simple share section in the smb.conf file looks similar to this:

```
[redbook]
comment = Redbook files
path = /redbook
browseable = yes
printable = no
writable = yes
write list = @users
```

Table 14 describes the most important parameters for creating a share.

Table 14. Share parameters

Parameter	Description
comment	This describes the function of the share.
admin users	This parameter is used to specify the users who have administrative privileges for the share. When they access the share they perform all operations as root.
path	Defines the full path to the directory you are sharing.
browseable	If this parameter is set to yes, you can see the share when you are browsing the resources on the Samba server. The value can be yes or no.
printable	This parameter specifies if the share is a print share. The value can be yes or no.
write list	Users specified in this list have write access to the share. If the name begins with @ it means a group name.
writable	This parameter specifies if the share is writable. The value can be yes or no.
read list	Users specified in this list have read access to the share. If the name begins with @ it means a group name.
read only	If this is set to yes, share is read only. The value can be yes or no.
valid users	This parameter specifies which users can access the share.

By using these parameters you can easily set up a new share. Each share definition starts with the share name in brackets “[]”. Below this name you can specify the values for the share parameters.

4.2.1.6 Share permissions

Although you can control the share permissions with share parameters, UNIX permissions are applied before the user can access files on the share. So you need to take care of UNIX permissions, so the user also has access to the shared directory under UNIX.

When a user creates a new file on the shared directory, the default create mask used is 0744. For directory creation, the default create mask is 0755. If you want, you can force a different creation mask. The parameters for doing this are explained in Table 15.

Table 15. Create mask parameters

Parameter	Description
create mask	This is used for file creation to mask against UNIX mask calculated from the DOS mode requested.
directory mask	This is used for directory creation to mask against UNIX mask calculated from the DOS mode requested.

4.2.1.7 Creating shares for home directories

For handling home directories Samba has a special share section called `[homes]`. This share definition is used for all home directories, so you do not need to create separate shares for each user.

When a client requests a connection to a file share, existing file shares are scanned. If a match is found, that share is used. If no match is found, the requested share is treated as a user name and validated by security. If the name exists and the password is correct, a share with that name is created by cloning the `[homes]` section. The home share definition uses the same parameters as a normal share definition. The following is an example of a home share definition in the `smb.conf` configuration file:

```
[homes]
comment = Home Directories
path = %H
valid users = %S
browseable = no
writable = yes
create mode = 0700
directory mode = 0700
```


As you can see, we used some variables in this definition, which are explained in Table 16.

Table 16. Variable description

Parameter	Description
%H	This variable represents the home directory of the user.
%S	The name of the current service, which is, in the case of home share, equal to username.

As you can see in the example, we used creation masks for the files and the directories in such a way that we forced all new files or directories to be accessible only by the owner of the home directory.

4.2.1.8 Creating a printer share

A Samba server uses the same procedure for printer shares as for the home shares. If all share definitions and user names are tested against the requested share name and the matched definition is still not found, Samba searches for a printer with that name (if the `[printers]` section exists). If the match is found in the printer definitions, that `[printers]` share section is cloned with the name of the requested service, which is really a printer name. The following is an example of the printers definition in the `smb.conf` configuration file:

```
[printers]
comment = All Printers
path = /var/spool/samba
browseable = no
# Set public = yes to allow user 'guest account' to print
guest ok = no
writable = no
printable = yes
create mask = 0700
```

As you can see, the `[printers]` section is just another share definition, when a user prints they basically copy the data into a spool directory, after that the data is handled by the local printing system. The only big difference between a printer share and other share definitions is that the `printable` parameter is set to “yes”. This means that a user can write a spool file to the directory specified under the share definition. If the share is printable, then it is also writable by default.

4.2.2 Starting and stopping the Samba server

You can start the Samba server by executing the command:

```
rcsmb start
```

You will see output similar to the following:

```
bash-2.04# rcsmb start
Starting SMB services:
```

The Samba server can be stopped by executing the command:

```
rcsmb stop
```

Whenever you make modifications to the `smb.conf` configuration file, you must restart the Samba server.

4.2.3 Starting Samba as startup service

You can configure your boot process so Samba is started at the boot. Edit the `/etc/rc.config` file and change `START_SMB="no"` to `START_SMB="yes"`. Make sure you run `SUSEconfig` to commit the changes.

4.2.4 Using SWAT

The Samba Web Administration Tool (SWAT) allows the remote configuration of the `smb.conf` configuration file through a Web browser. That means you can configure Samba in a GUI-like environment. SWAT itself is a small Web server and CGI scripting application, designed to run from `inetd`, provides access to the `smb.conf` configuration file.

An authorized user with the root password can configure the `smb.conf` configuration file via Web pages. SWAT also places help links to all configurable options on every page, which lets an administrator easily understand the effect of the changes.

SWAT is started with a TCP wrapper, so you can control who can access the SWAT service with the `/etc/hosts.deny` file. For example, if you want to access SWAT locally only, your `/etc/hosts.deny` file should look similar to this:

```
#
# hosts.deny    This file describes the names of the hosts which are
#              *not* allowed to use the local INET services, as decided
#              by the '/usr/sbin/tcpd' server.
#
# The portmap line is redundant, but it is left to remind you that
# the new secure portmap uses hosts.deny and hosts.allow.  In particular
# you should know that NFS uses portmap!
```

swat:ALL EXCEPT 127.0.0.1

To start SWAT point your favorite Web browser to the Internet address of your Samba server on port 901, as you can see in Figure 151.

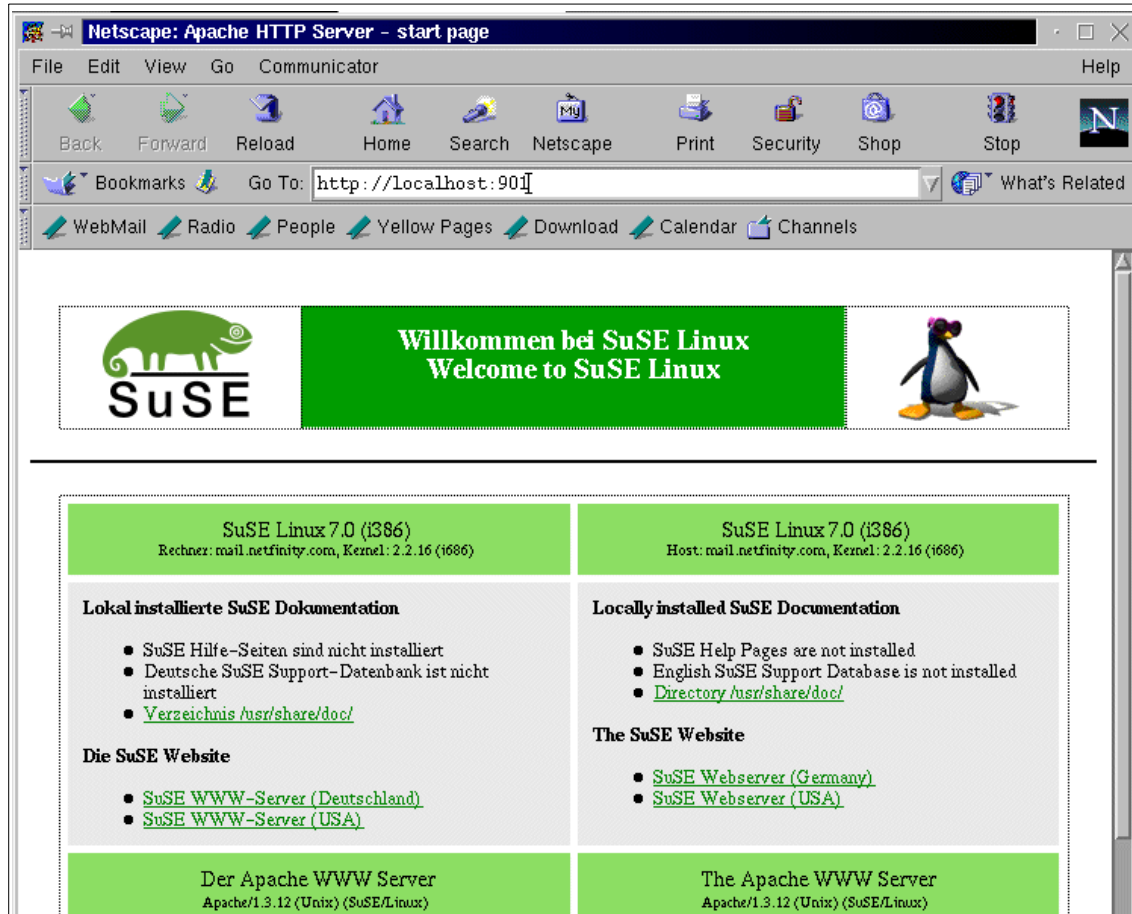


Figure 151. Starting SWAT

After you load the home page of SWAT, you will see a window similar to Figure 152.

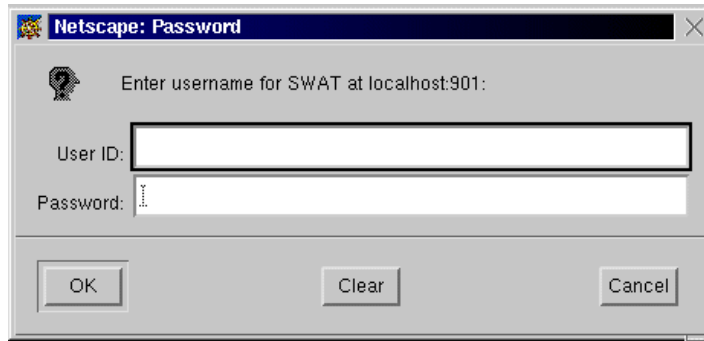


Figure 152. User authorization for SWAT

Type in the username and password of the Linux user defined on your Linux server. Click **OK** to continue. You will see a window similar to Figure 153.

Stop

Any Linux user can access SWAT, but only a root user can make changes.

Remember, when you are logging on to SWAT from a remote machine, you are sending passwords in plain text. This can be a security issue, so we recommend that you only administrate SWAT locally.

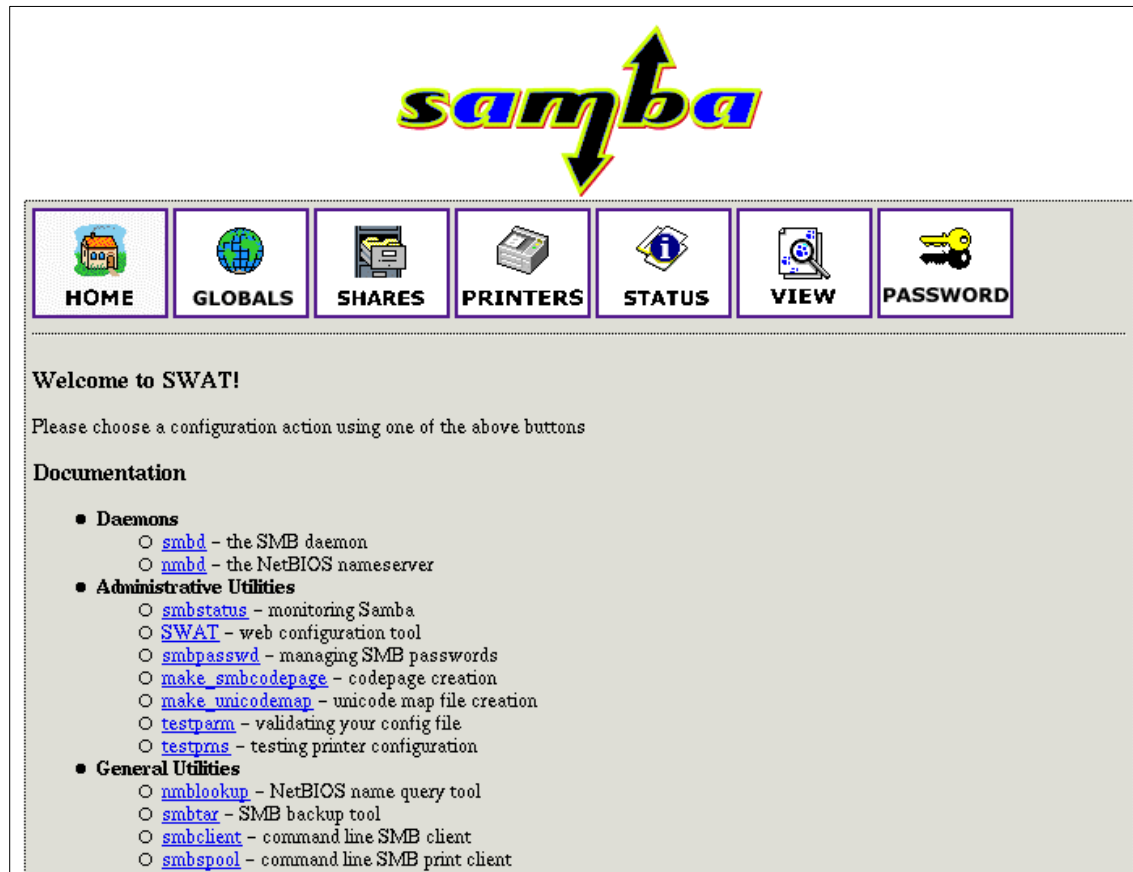


Figure 153. SWAT home page

As you can see in Figure 153, you have seven categories available:

1. Home - here you can find all the documentation you need about Samba.
2. Globals - here you can view and modify global parameters from the smb.conf configuration file.
3. Shares - here you can view, modify, and add shares.
4. Printers - here you can view, modify, and add printers.
5. Status - here you can check the current status of your Samba server.
6. View - here you can view current configuration of the smb.conf configuration file.
7. Passwords - here you can manage passwords for the Samba server.

Now we will briefly describe the functions available in SWAT.

Note

You can reach any of the seven functions on all SWAT Web pages. There are always icons for the functions on the top of each page.

After you make changes to smb.conf configuration file, the Samba server must be restarted.

4.2.4.1 Globals

When you click the **Globals** icon in the main SWAT window, and you will see a window similar to Figure 154.

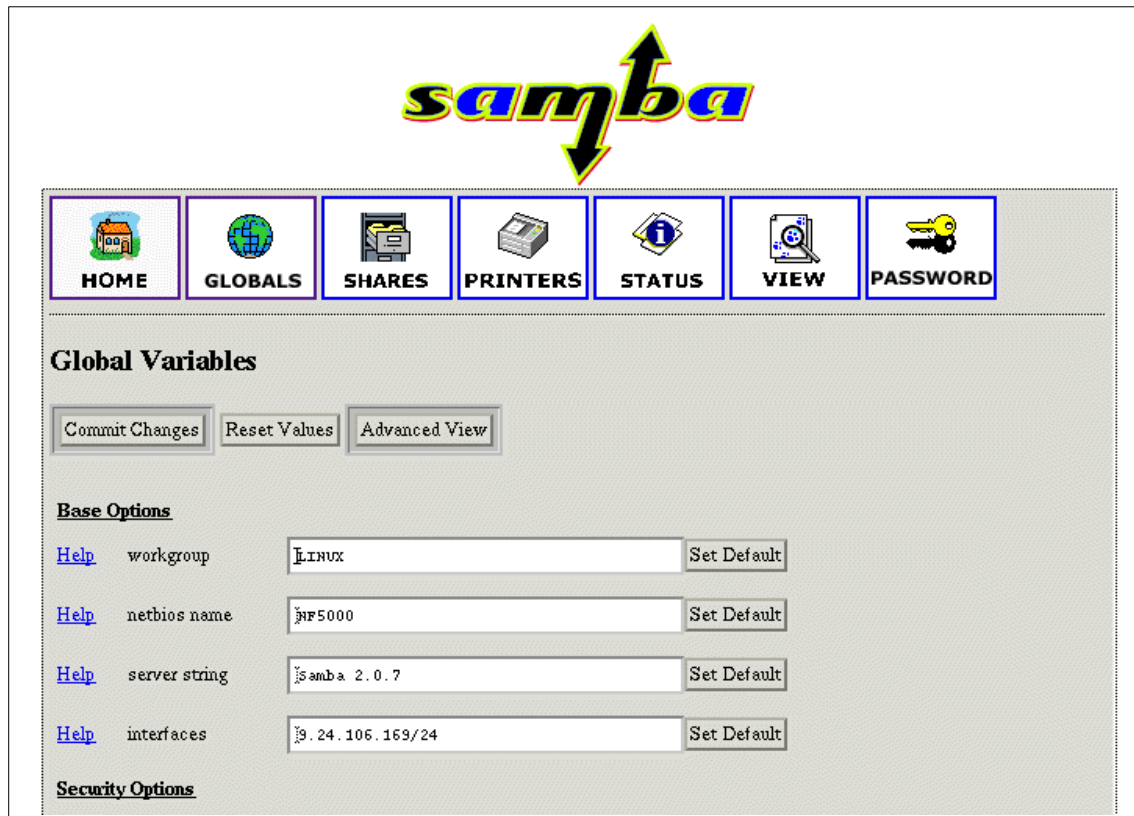


Figure 154. Global section in SWAT

In this window you can modify the global parameters for the Samba server. By default you will see the Basic View; if you want to see the Advanced View

click **Advanced View**. In the Advanced View you have all options available, while in the Basic View you can change only the basic options. To return from the Advanced View to the Basic View, click **Basic View**. After you have made your changes you can save them by clicking **Commit changes**.

4.2.4.2 Shares

When you click the **Shares** icon on any of the SWAT Web pages, you will see a window similar to Figure 155.

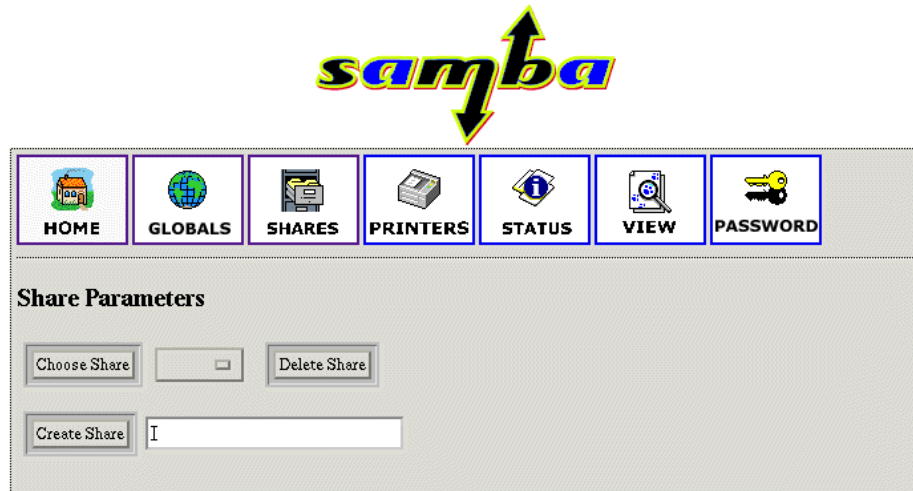


Figure 155. Shares section in SWAT

Here you can:

- View a defined share
- Delete shares
- Create a new share

4.2.4.3 Viewing or modifying an existing share

To view an already defined share, select the share from the field to the right of the **Choose Share** button, similar to Figure 156.

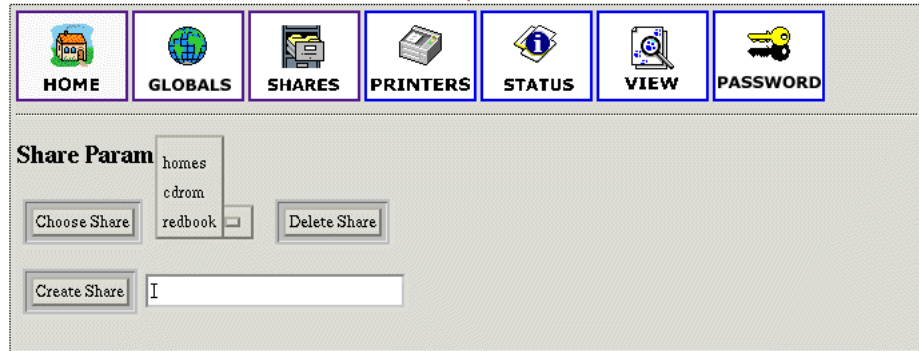


Figure 156. Choosing a share to view

After you have selected the share, click **Choose Share** to view the share properties. You will see a window similar to Figure 157.



HOME	GLOBALS	SHARES	PRINTERS	STATUS	VIEW	PASSWORD
------	---------	--------	----------	--------	------	----------

Share Parameters

redbook

Base Options

[Help](#) comment

[Help](#) path

Security Options

[Help](#) guest account

[Help](#) writeable

[Help](#) guest ok

[Help](#) hosts allow

[Help](#) hosts deny

Browse Options

[Help](#) browseable

Miscellaneous Options

[Help](#) available

Figure 157. Share properties

If you want to view all available parameters, click **Advanced View**. In this view you can also make changes and save them by clicking **Commit Changes**.

4.2.4.4 Deleting an existing share

To delete an existing share you must first select an already defined share similar to Figure 156. Then click **Delete Share**.

Important

The share will be deleted immediately and without warning.

After you have deleted the share, the Samba server must be restarted.

4.2.4.5 Creating a new share

To create a simple share, follow these steps:

1. Create a directory that will be used for the share. You can do this by executing this command from the terminal:

```
mkdir /home/public
```

In our example we created a “public” directory in the “home” directory.

2. Make sure that the UNIX permissions are set correctly in that directory, so that only intended users have access rights to it.
3. In the shares view of the SWAT Web pages, type in the name of the share you are creating, similar to Figure 158.

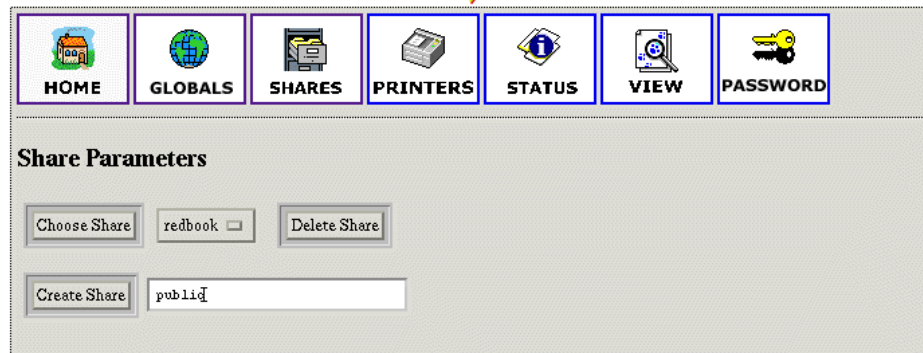


Figure 158. Entering the name for a new share

4. Click **Create Share** to continue, and you will see a window similar to Figure 159.



HOME	GLOBALS	SHARES	PRINTERS	STATUS	VIEW	PASSWORD
------	---------	--------	----------	--------	------	----------

Share Parameters

Choose Share: public Delete Share

Create Share:

Commit Changes Reset Values Advanced View

Base Options

[Help](#) comment Set Default

[Help](#) path Set Default

Security Options

[Help](#) guest account Set Default

[Help](#) writeable Set Default

[Help](#) guest ok Set Default

[Help](#) hosts allow Set Default

[Help](#) hosts deny Set Default

Browse Options

[Help](#) browseable Set Default

Miscellaneous Options

[Help](#) available Set Default

Figure 159. Entering the new share parameters

5. Fill in the relevant parameters. If you need to set more advanced parameters, click **Advanced View** and you will see all available parameters. After you have completed the configuration, click **Commit Changes** to save your new share.
6. You can see the changed smb.conf configuration file by selecting the **View** icon from the SWAT Web pages. You will see a window similar to Figure 160.

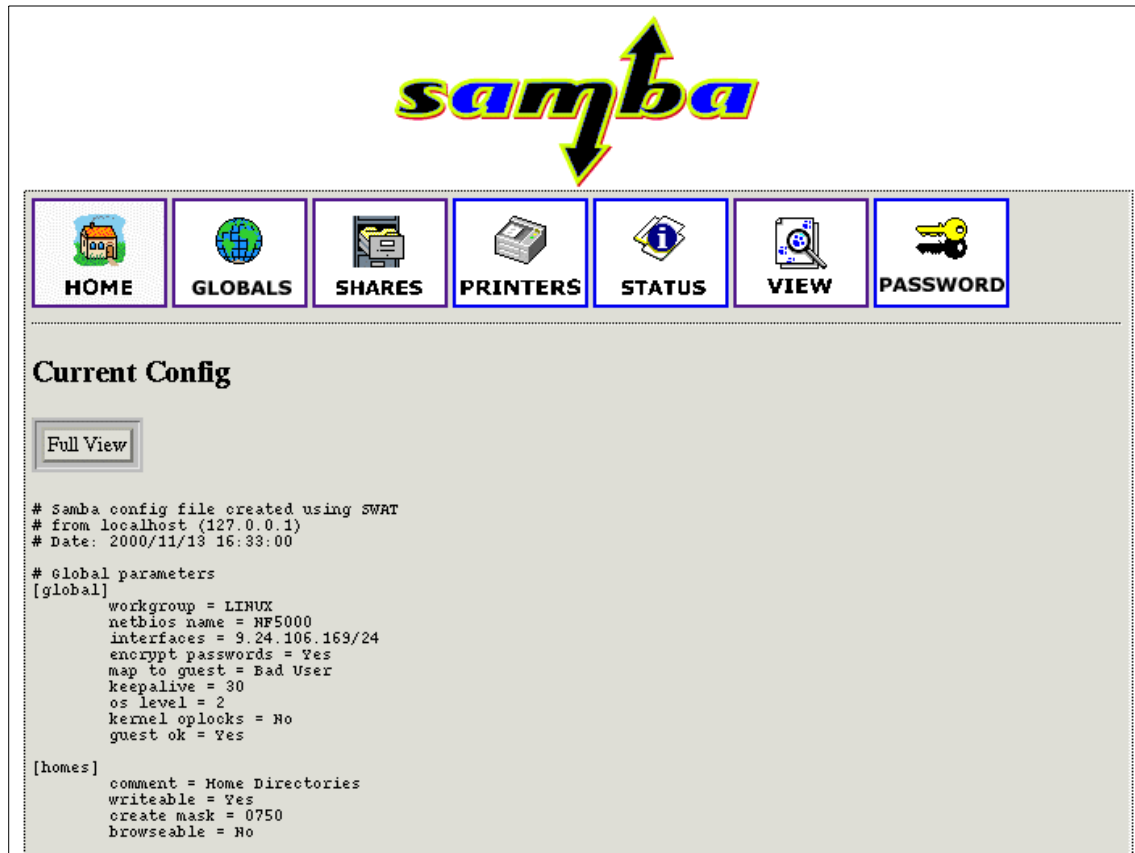


Figure 160. Viewing the smb.conf configuration file

7. Restart the Samba server.

Congratulations! You have just created your first usable share on the Samba server.

4.2.4.6 Restarting the Samba server

The Samba server can be restarted by clicking the **Status** icon on any of the SWAT configuration pages. You will see a window similar to Figure 161.

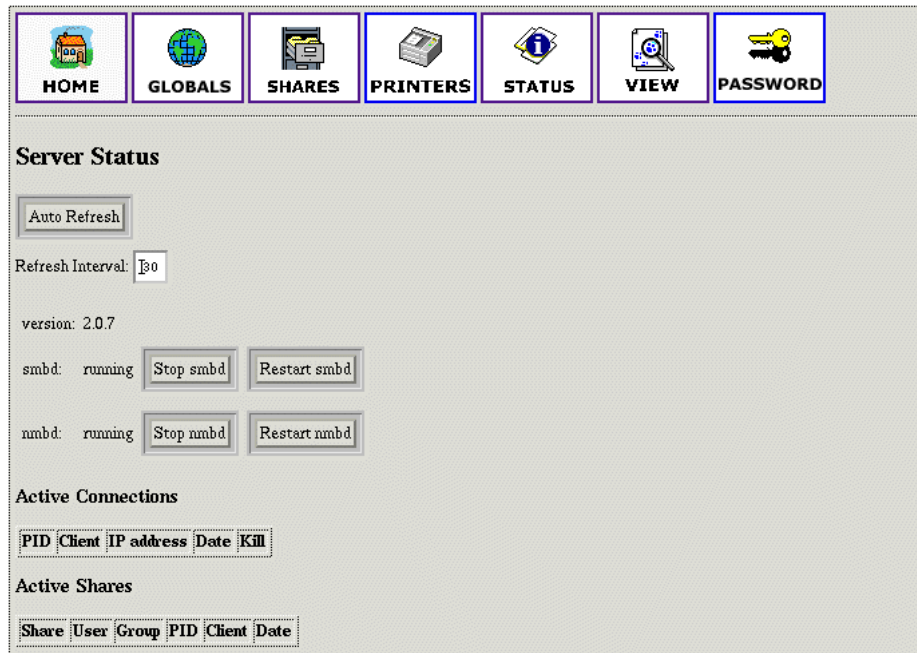


Figure 161. Restarting the Samba server

To restart the Samba server simply click **Restart smbd**. On this page you can also restart the WINS server by clicking **Restart nmbd**.

4.2.4.7 Printers

In the printer section you can view, modify, or add printers. The operations for handling printers are the same as for handling shares. You can access the printer settings by clicking the **Printers** icon on the SWAT Web page similar to Figure 162.

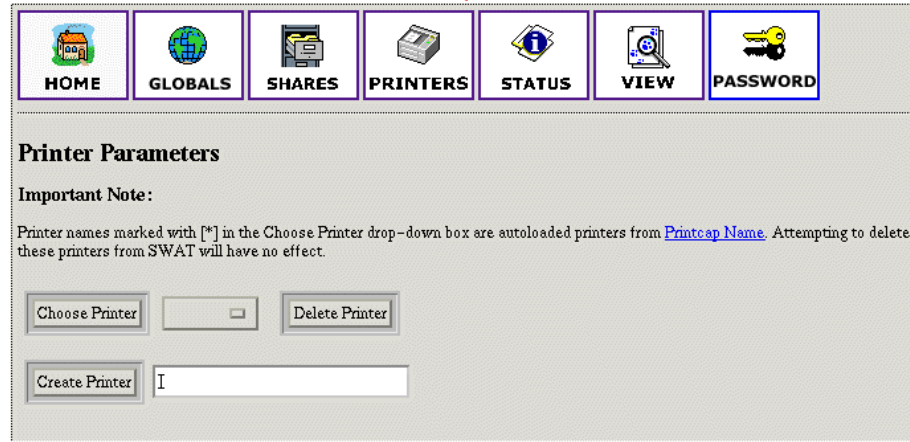


Figure 162. SWAT printers section

If you want to view the settings for a specific printer, select the printer from the list, as in Figure 163.

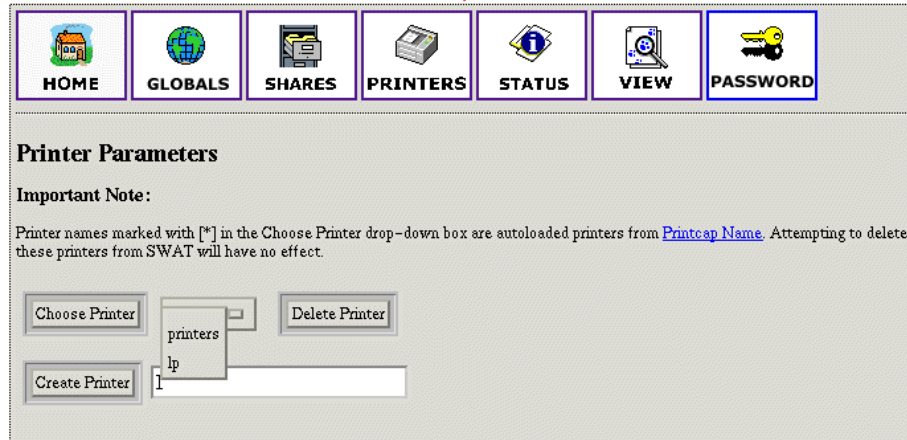


Figure 163. Selecting a printer

After you have selected the printer, click **Choose Printer** to view its properties. You will see a window similar to Figure 164.



HOME	GLOBALS	SHARES	PRINTERS	STATUS	VIEW	PASSWORD
------	---------	--------	----------	--------	------	----------

Printer Parameters

Important Note:
Printer names marked with [*] in the Choose Printer drop-down box are auto-loaded printers from [Printcap Name](#). Attempting to delete these printers from SWAT will have no effect.

Choose Printer: lp Delete Printer

Create Printer:

Commit Changes Reset Values Advanced View

Base Options

[Help](#) comment Set Default

[Help](#) path Set Default

Security Options

[Help](#) guest account Set Default

[Help](#) guest ok Set Default

[Help](#) hosts allow Set Default

[Help](#) hosts deny Set Default

Printing Options

[Help](#) printable Set Default

Browse Options

[Help](#) browseable Set Default

Miscellaneous Options

[Help](#) available Set Default

Figure 164. Printer properties

You can also modify the printer properties. When you are done, save the settings by clicking **Commit Changes**.

4.2.4.8 Status

In this section you can check the status of the Samba server. Here you can view all the connections and open files. You can also start or restart the Samba server or just its components.

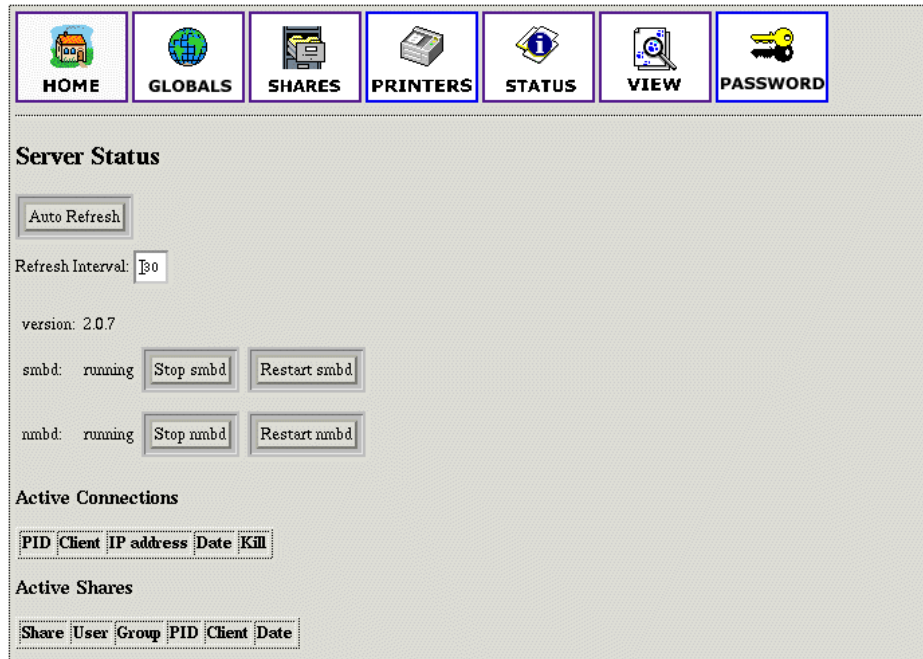


Figure 165. Status section

4.2.4.9 View

In this section you can view the current smb.conf configuration file. You can access printer settings by clicking the **View** icon on the SWAT Web pages similar to Figure 166.

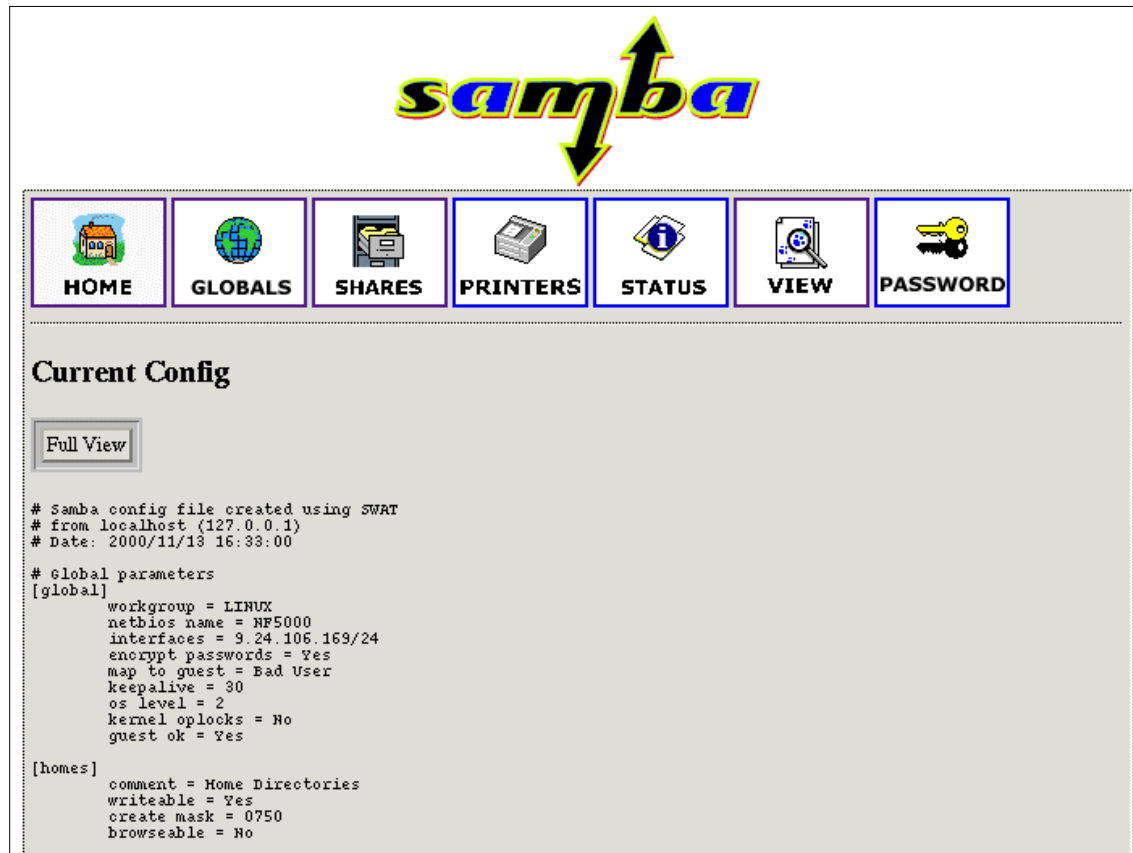
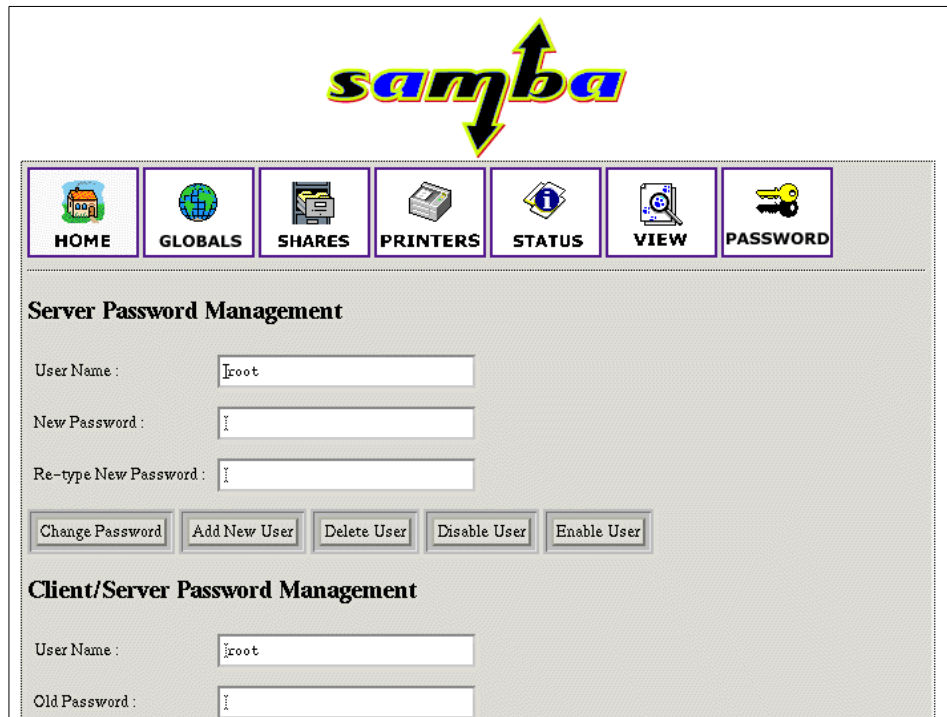


Figure 166. View section of SWAT

4.2.4.10 Password

In this section you can manage the passwords of all Samba users. You can access password settings by clicking the **Password** icon on the SWAT Web pages similar to Figure 167.



The screenshot displays the Samba web interface. At the top center is the Samba logo, which consists of the word "samba" in a stylized, colorful font with a double-headed arrow pointing up and down. Below the logo is a horizontal navigation bar with seven icons and labels: HOME (house icon), GLOBALS (globe icon), SHARES (folder icon), PRINTERS (printer icon), STATUS (info icon), VIEW (magnifying glass icon), and PASSWORD (key icon). The main content area is titled "Server Password Management" and contains three input fields: "User Name" (containing "root"), "New Password" (containing "1"), and "Re-type New Password" (containing "1"). Below these fields are five buttons: "Change Password", "Add New User", "Delete User", "Disable User", and "Enable User". The section below is titled "Client/Server Password Management" and contains two input fields: "User Name" (containing "root") and "Old Password" (containing "1").

Figure 167. Managing passwords

4.3 Sources and additional information

You can find more information on the official Samba project Web site at:

<http://www.samba.org>

There are always good how-to documents on the Linux Documentation project home page:

<http://www.linuxdoc.org/>

Chapter 5. DNS - Domain Name System

If you connect two or more computers to a network, they can share information and resources. However, these computers need to “talk in the same language” to be able to establish a connection. This “language” is called a network protocol. Today, the most popular communication protocol is TCP/IP. This is the protocol that is being used on the Internet and in many local area networks.

Hosts in a TCP/IP network communicate with each other by using unique IP addresses. These addresses consist of four 8-bit numbers (octets) that are divided by dots. For example, host A has the address 192.168.99.1, while host B uses 122.68.29.5.

However, this addressing scheme is not very comprehensible to human beings and it is almost impossible to memorize a number of hosts by their IP addresses. Therefore a naming scheme has been invented.

Each host has a host name (for example fred) and belongs to a certain domain (for example snake-oil.com). Domains can be organized in a hierarchical fashion and can consist of different subdomains (for example marketing.snake-oil.com). The combination of a host name and its domain name is called a fully qualified domain name (FQDN) (for example fred.marketing.snake-oil.com). Since domains are hierarchical, it is possible to have more hosts with the same host name in different subdomains. Therefore, fred.marketing.snake-oil.com can be a different host from fred.management.snake-oil.com. If you want these hosts to be addressable from the Internet, you need to register your domain name with a central registry. There are several top-level domains, such as .com, .org or .net. In addition to these generic top-level domains, each country in the world has its own country code as the top-level domain. For example, Germany has .de, Denmark has .dk, and Finland uses .fi.

Since the hosts internally still use their IP addresses to communicate, there needs to be a mapping between host names and the corresponding IP address. There are two ways this can be implemented.

All host names of a network, including their IP addresses, are put into a static text file. This file has to be copied on each host that wants to communicate with the others by name. As soon as a host has been added or removed from the network, or an IP address or host name has changed, and the host files on all computers have to be adjusted accordingly. This can get very tedious, if the number of hosts is large.

This is where the Domain Name System (DNS) comes in. The following description of DNS is very simplified, but it should give you a rough picture of what DNS is all about.

Instead of maintaining a separate host file on each machine, there is a central server that carries a list of all hosts and IP addresses of its domain. All clients now send their host name resolution request to this central server instead of looking in a local table. The name server will look up the requested host name and return the respective IP address. The opposite is also possible: the client can also ask for a host name that belongs to a certain IP address. If a client asks for an IP address of another domain, the local domain name server will forward the request to the next name server above in its hierarchy, if it cannot answer the request by itself. Therefore changes to the table of host names have to be made at one central point only rather than on all participants of the network.

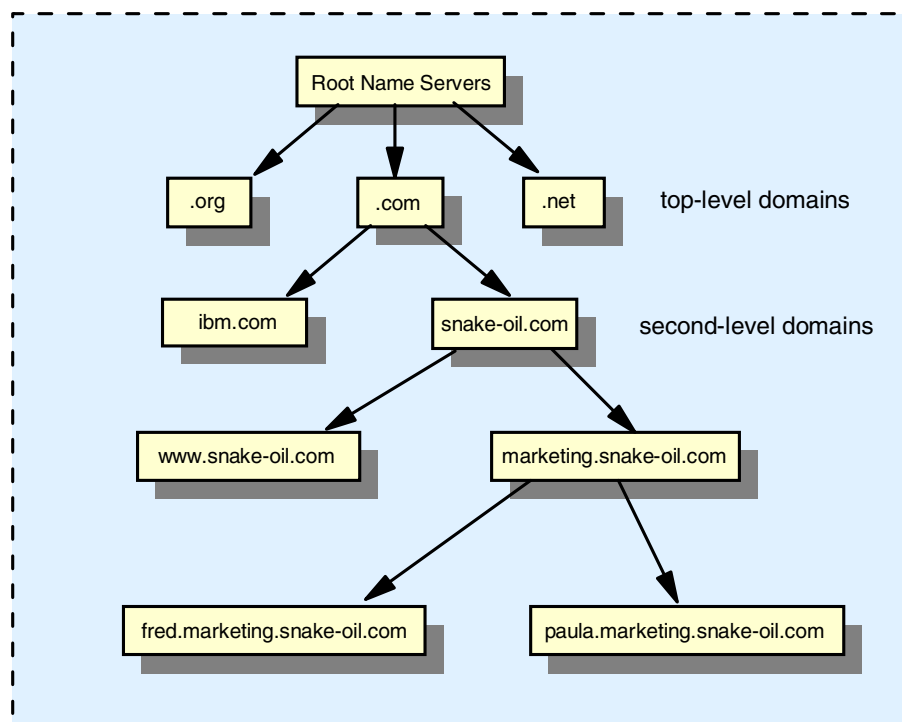


Figure 168. Internet domain hierarchy

This chapter describes how to set up a name server for a local domain and how to maintain a host list for this domain.

5.1 Installation of software

The server that will be the DNS server needs to have a working TCP/IP network connection to the other hosts in its network before we start. The program that is responsible for this service is called *named* and belongs to the software package *bind*, which is maintained by Paul Vixie for The Internet Software Consortium. There are two major versions of *bind*: *bind4* and *bind8*. We will focus on the new version *bind8*, because it is more secure and is designed to replace *bind4* in the future. Most Linux distributions already contain a precompiled and preconfigured package for *bind8*.

Note

The package *bind8* has been split up into two separate packages in SuSE Linux 7.0: *bind8*, which contains the actual server program, and *bindutil*, which contains the utilities such as *nslookup*, *dig* and *host*. We recommend that you install both on the server. A client machine only needs the *bindutil* package.

Make sure that the *bind* package is actually installed. In SuSE Linux, you can use the RPM package manager to query the database of installed packages by entering the following command:

```
rpm -q bind8
```

If the package is already installed, RPM will return the version and build number of this package:

```
bind8-8.2.3-34
```

If it is not installed, you will receive the following message:

```
package bind8 is not installed.
```

You will then have to install this package first. Please refer to 3.1, “Adding and removing software packages using YaST ’ on page 111 for how to install software packages. The package *bind8* is located in series *n - Network-Support (TCP/IP, UUCP, Mail, News)*. Quit YaST to return to the command line after installing the package.

5.2 DNS sample configuration

Configuring DNS can be very complex, depending on the intended functionality. Covering this in depth is beyond the scope of this chapter. We

will therefore focus on very a simplified example and recommend that you take a look at the very informative DNS how-to at:

<http://www.linuxdoc.org/HOWTO/DNS-HOWTO.html>

or at `/usr/share/doc/howto/en/DNS-HOWTO.gz` on your local file system for further information on DNS and bind.

We will construct a simple example: The company Snake Oil Ltd. wants to set up a local DNS server for their internal network (the internal IP address range is 192.168.99.xxx/24, a Class C network). They chose snake-oil.com as their local domain name. The network is also connected to the Internet. The name server will be configured to answer all requests about the local (internal) snake-oil.com domain and forward all other requests to the ISP's name server (ns.bigisp.com, fictional IP address 155.3.12.1) as a caching name server.

We begin with a simple example. At first the local DNS will be configured to act as a caching-only name server. This means that it forwards all requests to the ISP's name server(s) (forwarders) and caches all answers for further requests from its clients. This reduces the network traffic on the outside line.

Put the following lines in the `/etc/resolv.conf` file:

```
search snake-oil.com
nameserver 127.0.0.1
```

This will make sure that the server itself will use its local name server for host name resolution.

In SuSE Linux, you can use YaST to modify this entry. Choose **System administration -> Network configuration -> Configuration nameserver**. Enter the IP address 127.0.0.1 and your domain. To enter this dialog directly from the command line, enter the following command:

```
yast --mask nameserver --autoexit
```

The name server's main configuration file is `/etc/named.conf`. Most distributions ship with a very detailed example configuration file; you might want to save this for future reference. We will create a new file from scratch. Open up a text editor and create a new `/etc/named.conf` according to Figure 169:


```

options {
    directory "/var/named";
    pid-file "/var/named/slave/named.pid";
    listen-on { any; };
    forward only;
    forwarders { 155.3.12.1; };
    sortlist {
        { localhost; localnets; };
        { localnets; };
    };
};

logging {
    category lame-servers { null; };
    category cname { null; };
};

zone "localhost" IN {
    type master;
    file "localhost.zone";
    check-names fail;
    allow-update { none; };
};

zone "0.0.127.in-addr.arpa" IN {
    type master;
    file "127.0.0.zone";
    check-names fail;
    allow-update { none; };
};

```

Figure 169. *Named.conf* text file

Replace the IP address in the `forwarders` field with your ISP's name server IP address.

You also need to create the following `/var/named/localhost.zone` file:

```

$ORIGIN localhost.
@           1D IN SOA      @ root (
                    42           ; serial (d. adams)
                    3H           ; refresh
                    15M          ; retry
                    1W           ; expiry
                    1D )         ; minimum

                    1D IN NS      @
                    1D IN A      127.0.0.1

```

Figure 170. *localhost.zone* text file

Create the file `/var/named/127.0.0.zone` with the following content:

```

$ORIGIN 0.0.127.in-addr.arpa.

@           1D IN SOA      localhost. root.localhost. (
                    42           ; serial (d. adams)
                    3H           ; refresh
                    15M          ; retry
                    1W           ; expiry
                    1D )         ; minimum

1           1D IN NS      localhost.
1           1D IN PTR     localhost.

```

Figure 171. The 127.0.0. zone text file

Your network clients should all be configured to query the local DNS server's IP address instead of your ISP's name server.

You can now start the server with the command:

```
rcnamed start
```

Check `/var/log/messages` for the startup messages. The name server should now resolve DNS queries from its clients by forwarding them to the ISP's name server. You can verify this with the commands `host <somehostname>` and `nslookup`.

If you want the name server to be started at the next system reboot, set the variable `START_NAMED` in `/etc/rc.config` to "yes". See 3.6, "Changing the configuration file with YaST" on page 132 for how to do this.

In the following step, we will configure the server to act as a primary name server for the local domain `snake-oil.com`. Stop the name server with `rcnamed stop` and edit the file `/etc/named.conf` so that it looks like Figure 172:

```

options {
    directory "/var/named";
    pid-file "/var/named/slave/named.pid";
    listen-on { any; };
    forward only;
    forwarders {9.24.106.15;};
    sortlist {
        { localhost; localnets; };
        { localnets; };
    };
};

logging {
    category lame-servers { null; };
    category cname { null; };
};

zone "." {
    type hint;
    file "root.hint";
};

zone "localhost" IN {
    type master;
    file "localhost.zone";
    check-names fail;
    allow-update { none; };
};

zone "0.0.127.in-addr.arpa" IN {
    type master;
    file "127.0.0.zone";
    check-names fail;
    allow-update { none; };
};

zone "snake-oil.com" {
    type master;
    file "snake-oil.zone";
};

zone "99.168.192.IN-ADDR.APRA" {
    type master;
    file "snake-oil.rev";
};

```

Figure 172. Name.conf text file

We have now added the zone files (the databases) needed for our local domain "snake-oil.com". The file /var/named/snake-oil.zone (Figure 173) is responsible for the mapping of host names to IP addresses.

```

; Zone file for snake-oil.com
;
@      IN      SOA      ns.snake-oil.com. hostmaster.snake-oil.com. (
                        199910011      ; serial, todays date + todays serial #
                        8H              ; refresh, seconds
                        2H              ; retry, seconds
                        1W              ; expire, seconds
                        1D )            ; minimum, seconds
;
;
;      NS      ns              ; Inet Address of name server
;      MX      10 mail         ; Primary Mail Exchanger
;      MX      20 mail.bigisp.com. ; Secondary Mail Exchanger
;
localhost      A      127.0.0.1
gw              A      192.168.99.1
ns              A      192.168.99.2
fred           A      192.168.99.3
mail           A      192.168.99.4
ftp            A      192.168.99.5
www            A      192.168.99.6

```

Figure 173. Snake-oil.com text file

You should also create the zone file `/var/named/snake-oil.rev`. This is necessary for reverse name lookups, for example, if you need to resolve an IP address to its host name.

The MX record in the zone file tells other hosts on the Internet what mail server services this domain. In our case, mail for `fred@snake-oil.com` will be relayed through `mail.snake-oil.com`, and as a backup, through `mail.bigisp.com`. The 10 and 20 in the second column signifies the priority of the mail servers, effectively providing redundancy.

```

@      IN      SOA      ns.snake-oil.com. hostmaster.snake-oli.com. (
                        199910011 ; Serial, todays date + todays serial
                        8H      ; Refresh
                        2H      ; Retry
                        1W      ; Expire
                        1D)     ; Minimum TTL
;
;      NS      ns.snake-oil.com.
;
1      PTR     gw.snake-oil.com.
2      PTR     ns.snake-oil.com.
3      PTR     fred.snake-oil.com.
4      PTR     mail.snake-oil.com.
5      PTR     ftp.snake-oil.com.
6      PTR     www.snake-oil.com.
.

```

Figure 174. Snake-oil.rev zone file

Now let the name server reload its configuration again by running `rndc restart`. Have a look at the messages in `/var/log/messages`. If everything went well, you should see messages similar to the following:

```
Oct 26 18:03:20 ns named[14870]: starting
Oct 26 18:03:20 ns named[14870]: cache zone "" (IN) loaded (serial 0)
Oct 26 18:03:20 ns named[14870]: master zone "localhost" (IN) loaded (serial 42)
Oct 26 18:03:20 ns named[14870]: master zone "0.0.127.in-addr.arpa" (IN) loaded serial 42)
Oct 26 18:03:20 ns named[14870]: master zone "snake-oil.com" (IN) loaded (serial 199910011)
Oct 26 18:03:20 ns named[14870]: master zone "99.168.192.IN-ADDR.APRA" (IN) load ed (se ial 199910011)
Oct 26 18:03:20 ns named[14870]: listening on [127.0.0.1].53 (lo)
Oct 26 18:03:20 ns named[14870]: listening on [9.24.105.210].53 (eth0)
Oct 26 18:03:20 ns named[14870]: Forwarding source address is [0.0.0.0].1041
Oct 26 18:03:20 ns named[14871]: Ready to answer queries.
```

Figure 175. `/var/log/messages` file

Your name server should now correctly resolve host names for the snake-oil domain as well.

5.3 Configuration tips

Use the `listen-on` directive in the options section of the `named.conf` file. For each interface a name server listens on, a pair of filehandles is opened. On a busy name server, saving every filehandle is a big win.

Check the `/var/log/messages` file from time to time for errors. Named is pretty verbose in its error messages.

If you are constantly adding, removing or just making modifications to your zone records, you might want to have a look at the `nsupdate` tool, which also belongs to the `bind8` package.

Chapter 6. Secure Shell

Security is a big issue in networks today, and as such many tools have been developed to make securing a network just that little bit easier.

The most widely used application in communicating and maintaining machines has been the Secure Shell (SSH). SSH provides a way to encrypt every aspect of a connection between two computers based on public and private key technology. Passwords are not sent as clear text as is common with the “normal” network services, but a challenge is sent by the server on connection that the client must answer, all encrypted.

Once a connection has been established, all communications on the SSH link are encrypted based on the public and private key pair.

6.1 Installing SSH

The US version of SuSE Linux does not include SSH by default, due to US export restrictions on strong cryptography."

You can download SSH for SuSE Linux at:

```
ftp://ftp.gwdg.de/pub/linux/suse/7.0/i386.de/suse/sec1/openssh.rpm
```

Install SSH by typing:

```
rpm -i openssh.rpm
```

You will notice that it will update `/etc/rc.config`, so you will have to run `SuSEconfig` to update your system.

6.2 Configuring SSH

To set up a secure system, we need to tell the SSH daemon not to allow a fallback to password authentication, as it still allows “guessing” of users’ passwords. It is also advisable to restrict access to the SSH service to a certain network device. This is not needed if you wish everyone to be able to access the system based on public/private key authenticating.

Open the file `/etc/ssh/sshd_config` and change **PasswordAuthentication yes** to **PasswordAuthentication no**. This disables users to log in using a password and forces a key challenge.

To change which network device SSH will listen on, add the IP addresses of the network adapters to the **ListenAddress** parameter. Using an IP address of 0.0.0.0 will tell the SSH daemon to listen on all network devices in the system.

6.2.1 Host key generation

Start the SSH daemon by typing:

```
rcsshd start
```

The SSH daemon will create a random host key to uniquely identify the system.

```
Generating /etc/ssh/ssh_host_key.  
Generating RSA keys: .....0000000.....0000000  
Key generation complete.  
Your identification has been saved in /etc/ssh/ssh_host_key.  
Your public key has been saved in /etc/ssh/ssh_host_key.pub.  
The key fingerprint is:  
5c:6b:9b:00:57:ff:dd4b:57:6c:43:12:76:87:c9:57 root@mail  
Generating /etc/ssh/ssh_host_dsa_key.  
Generating DSA parameter and key. It can take a long time.  
Your identification has been saved in /etc/ssh/ssh_host_dsa_key.  
Your public key has been saved in /etc/ssh/ssh_host_dsa_key.pub.  
The key fingerprint is:  
d5:7b:b8:6a:6c:de:16:d1:71:84:90:51:fe:59:41:40 root@mail  
Starting SSH daemon:  
mail:/floppy #
```

Figure 176. Starting the SSH daemon for the first time

6.2.2 User key generation

Once this has been done we need to create a public/private key pair for users. To do this you must run

```
ssh-keygen
```

at a command prompt on the client you wish to connect from. You will be asked where to store your key pair. The default values are fine for this.


```
mail:/etc/ssh # ssh-keygen
Generating RSA keys: .....0000000.....0000000
Key generation complete.
Enter file in which to save the key (/root/.ssh/identify):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/identify.
Your public key has been saved in /root/.ssh/identify.pub.
The key fingerprint is:
6c:4d:e3:71:08:0f:53:36:33:a1:d5:4e:aa:58:7b:d6 root@mail
mail:/etc/ssh #
```

Figure 177. Creating a key pair

You will be asked for a pass phrase. This allows a higher degree of protection from malicious users seeing your keys. The key pair will be encrypted using your pass phrase, and you will be asked for your pass phrase each time you log in to the system. You can use sentences for your pass phrase; the longer it is, the harder it is to crack. And as a general warning do not use a common saying like:

“The quick brown fox jumped over the lazy dog”

You would be surprised how common it is.

6.2.3 Configuring connections

Once your key pair has been created, you will need to allow the keys to be used to log in to the system. This can be done by copying your public key to the file `authorized_keys` in the `~/.ssh` directory (the `.ssh` directory in your home directory) on the server you wish to connect to.

```
cat identity.pub >> authorized_keys
```

Using `>>` concatenates the `identity.pub` file to the `authorized_keys` file. This way you can keep on adding public keys if you are connecting from more than one host.

You need to use your private key to log in to the server. Your private key is your pass to the server. If anyone has your private key, they have access to your files, so do not lose it!

You can now log in to the server using the SSH client:

```
bash-2.04# ssh 192.168.158.128 -v
SSH Version OpenSSH_2.1.1, protocol versions 1.5/2.0.
Compiled with SSL (0x0090581f).
debug: Reading configuration data /etc/ssh/ssh_config
debug: Applying options for *
debug: Seeding random number generator
debug: ssh_connect: getuid 0 geteuid 0 anon 0
debug: Connecting to 192.168.158.128 [192.168.158.128] port 22.
debug: Seeding random number generator
debug: Allocated local port 678.
debug: Connection established.
debug: Remote protocol version 1.99, remote software version OpenSSH_2.1.1
debug: Local version string SSH-1.5-OpenSSH_2.1.1
debug: Waiting for server public key.
debug: Received server public key (768 bits) and host key (1024 bits).
debug: Host '192.168.158.128' is known and matches the RSA host key.
debug: Seeding random number generator
debug: Encryption type: 3des
debug: Sent encrypted session key.
debug: Installing crc compensation attack detector.
debug: Received encrypted confirmation.
debug: Trying RSA authentication with key 'root@mail'
debug: Received RSA challenge from server.
debug: Sending response to host key RSA challenge.
debug: Remote: RSA authentication accepted.
debug: RSA authentication accepted by server.
debug: Requesting pty.
debug: Requesting shell.
debug: Entering interactive session.
Last login: Mon Nov 13 17:51:21 2000 from 192.168.158.128
Have a lot of fun...
mail:~ #
```

Figure 178. Verbose SSH login

You now have a secure, encrypted way for you and your users to log in to your system. SSH can be used to tunnel connections for many services, providing a secure environment to work in. For more information about this, take a look at the SSH man page.

Chapter 7. Apache and IBM HTTP Servers

The Apache Web server is the most popular Web server software on today's Internet. According to the NetCraft Web server survey at <http://www.netcraft.com/survey/>, approximately 60% of all surveyed Web servers (more than 13 million) were running a version of Apache (as of the time of this writing). Apache is a very successful collaborative Open Source project. The Web site for Apache is <http://www.apache.org>. Because of the free availability of the full source code, it is a very flexible and powerful Web server solution. There are also a lot of additional modules, which can be used in combination with the Apache main program. Some popular examples are PHP (PHP: Hypertext Preprocessor, an embedded HTML scripting language), `mod_perl` (an embedded perl interpreter) and `mod_ssl` for secure transactions. More Apache modules can be downloaded from the Apache Module Registry at:

<http://modules.apache.org>.

Some of key features of Apache are:

- Implements the latest protocols, including HTTP/1.1 (RFC2068).
- Is highly configurable and extensible with third-party modules.
- Can be customized by writing “modules” using the Apache module API.
- Provides full source code and comes with an unrestrictive license.
- Runs on most versions of UNIX (including Linux) without modification.
- DBM databases for authentication, which allow you to easily set up password-protected pages with enormous numbers of authorized users without bogging down the server. A wide variety of SQL databases can be used for authentication too (using additional modules).
- Customized responses to errors and problems, which allow you to set up files, or even CGI scripts, which are returned by the server in response to errors and problems. For example, you can set up a script to intercept 500 server errors and perform on-the-fly diagnostics for both users and yourself.
- Multiple `DirectoryIndex` directives, which allow you to “say” `DirectoryIndex index.html index.cgi`, which instructs the server to either send back `index.html` or run `index.cgi` when a directory URL is requested, whichever it finds in the directory.
- Unlimited numbers of aliases and redirect directives that may be declared in the config files.

- Content negotiation, the ability to automatically serve clients of varying sophistication and HTML level compliance, with documents that offer the best representation of information that the client is capable of accepting.
- Multi-homed servers, which allow the server to distinguish between requests made to different IP addresses (mapped to the same machine).

7.1 The IBM HTTP Server

The IBM HTTP Server powered by Apache is based on the Apache HTTP Server. In addition to Linux, this HTTP Server also runs on AIX, Solaris and Windows NT. See the home page at:

<http://www.ibm.com/software/web servers/httpservers/>

IBM HTTP Server for Linux offers the following additional features:

- Remote Configuration: a browser-based configuration tool to allow manipulation of the server configuration via a GUI.
- SNMP Support: Simple Network Management Protocol (SNMP) is a well-established protocol for managing and gathering information about servers remotely. This new support allows IBM HTTP Server to be managed by the SNMP protocol.
- LDAP: The IBM HTTP Server Lightweight Directory Access Protocol (LDAP) plug-in allows authentication and authorization (which is required when accessing a protected resource) to be performed by an LDAP server, thereby greatly decreasing the administrative overhead for maintaining user and group information locally for each Web server.
- Machine Translation Support: This new function, when used with an available IBM Machine Translation Engine, enables the IBM HTTP Server to translate English Web pages into other languages without human intervention. This permits Web site visitors to read the page in their native language, effectively broadening the reach of your Web site. IBM Machine Translation Engines are included in the WebSphere Application Server 3.0 and include German, Simplified Chinese and Traditional Chinese. Additional languages will be available in the future.
- Support for SSL secure connections: The IBM HTTP Server powered by Apache supports both the SSL Version 2 and SSL Version 3 protocols. This protocol, implemented using IBM security libraries, ensures that data transferred between a client and a server remains private. Once your server has a digital certificate, SSL-enabled browsers such as Netscape Navigator and Microsoft Internet Explorer can communicate securely with your server using the SSL protocol. The IBM HTTP Server powered by

Apache supports client authentication, configurable cipher specifications, and session ID caching for improving SSL performance on the UNIX platforms.

- **Fast Response Cache Accelerator:** The Cache Accelerator can dramatically improve the performance of the IBM HTTP Server powered by Apache when serving static pages, for example, text and image files. Because the Cache Accelerator cache is automatically loaded during server operation, you are not required to list the files to be cached in your server configuration file. In addition, the server will automatically recache changed pages and remove outdated pages from the cache. The Cache Accelerator provides support for caching on Web servers with single and multiple TCP/IP adapters.

7.2 Apache HTTP Server installation

The Apache HTTP Server is installed and started by default on SuSE Linux, because it is used for the online help system. You can verify the installation by querying the RPM database:

```
rpm -q apache
```

This command will return either the version number of the installed package or an error message, if the package is not installed. Refer to 3.1, “Adding and removing software packages using YaST” on page 111 for how to install the package if it is missing. The package window is located in series n - Network-Support (TCP/IP, UUCP, Mail, News). Apache will be automatically started on bootup, if the variable `START_HTTP` in the central configuration file `/etc/rc.config` is set to `yes`. See 3.6, “Changing the configuration file with YaST” on page 132 for methods to modify this variable. To start, stop or reload the server (after a configuration change), run the script:



```
/usr/sbin/rcapache (start|stop|reload).
```

This file is a symbolic link to the init script in:

```
/sbin/init.d/apache.
```

In the SuSE default installation, Apache will serve HTML documents from the directory `/usr/local/httpd/htdocs` and CGI scripts from `/usr/local/httpd/cgi-bin`. If you installed the PHP module (`mod_php`), it will also execute PHP code, if the file ends in `.php3`. The access log file is in `/var/log/httpd.access_log` the error log file is `/var/log/httpd.error_log`. The Apache configuration files reside in the subdirectory `/etc/httpd`.

If you now point your browser to the server's IP address, you should see the following start page (/usr/local/httpd/htdocs/index.html) when the Apache HTTP Server is running:

 <div style="text-align: center; background-color: #008000; color: white; padding: 5px;"> Willkommen bei SuSE Linux Welcome to SuSE Linux </div> 	
SuSE Linux 7.0 (i386) <small>Rechner: netfinity.ibm.com, Kernel: 2.2.16 (i686)</small>	SuSE Linux 7.0 (i386) <small>Host: netfinity.ibm.com, Kernel: 2.2.16 (i686)</small>
Lokal installierte SuSE Dokumentation <ul style="list-style-type: none"> • SuSE Hilfe-Seiten • Deutsche SuSE Support-Datenbank ist nicht installiert • Verzeichnis /usr/share/doc/ Die SuSE Website <ul style="list-style-type: none"> • SuSE WWW-Server (Deutschland) • SuSE WWW-Server (USA) 	Locally installed SuSE Documentation <ul style="list-style-type: none"> • SuSE Help Pages • SuSE Support Database • Directory /usr/share/doc/ The SuSE Website <ul style="list-style-type: none"> • SuSE Webserver (Germany) • SuSE Webserver (USA)
Der Apache WWW Server <small>Apache/1.3.12 (Unix) (SuSE/Linux)</small>	The Apache WWW Server <small>Apache/1.3.12 (Unix) (SuSE/Linux)</small>
Dokumentation <ul style="list-style-type: none"> • Dokumentation zum Apache-Server (Englisch) • PHP ist nicht installiert Test- und Beispiel-Seiten <ul style="list-style-type: none"> • Apache perl Module (mod_perl) ist nicht installiert • Apache PHP3 Module (mod_php) ist nicht installiert 	Documentation <ul style="list-style-type: none"> • Apache webserver documentation • PHP is not installed Test and Example Pages <ul style="list-style-type: none"> • Apache perl module (mod_perl) is not installed • Apache PHP3 Module (mod_php) is not installed



SuSE GmbH Nürnberg, 2000; <feedback@suse.de>

Figure 179. Apache startup page on SuSE Linux

7.3 IBM HTTP Server installation

To install the IBM HTTP Server on SuSE Linux, you need to perform the following steps.

Before you are able to download the server files, you will have to register with IBM. It only takes a few minutes, and it allows you to quickly log in to get applications and documentation, etc.

For the IBM HTTP Server and the remote administration capabilities, download the .tar file from the Web page:

<http://www.ibm.com/software/webservers/httpservers/download.html>

The HTTPServer_linux_128_tar.tar.gz (or HTTPServer_linux_56_tar.tar.gz for 56-bit encryption) file contains the following packages:

- IBM_HTTP_Server-1.3.12-0.i386.rpm - IBM HTTP Server
- IBM_Apache_Source-1.3.12-0.i386.rpm - Apache 1.3.12 source
- IBM_Admin_Server-1.3.12-0.i386.rpm - Administration Server
- IBM_Admin_Server_Forms-1.3.12-0.i386.rpm - Administration Server Web forms
- gsk4bas-4.0-3.57.i386.rpm - Security library
- IBM_SSL_128-1.3.12-0.i386.rpm - 128-bit SSL library
- or
- IBM_SSL_56-1.3.12-0.i386.rpm - 56-bit SSL library
- IBM_SSL_Base-1.3.12-0.i386.rpm - SSL module
- IBM_Machine_Translation-1.3.12-0.i386.rpm - Gateway to IBM MT engine)
- IBM_SNMP-1.3.12-0.i386.rpm - SNMP client

We will not be discussing installation of the SSL and SNMP modules here. For more information about these, read the documentation included in the server by clicking **View Documentation** on the start page of the server site.

After you have downloaded the “gzipped tarball”, move it to the /tmp directory and extract it with the command:

```
tar -zxvf HTTPServer_linux_128_tar.tar.gz
```

This will extract the RPM files listed above from the tar archive into the subdirectory /tmp/IHS-1.3.12. You now need to become the root user (if you

are not already). To avoid resource conflicts, you first have to shut down the currently running Apache Web server (if installed), by executing the following command:

```
rcapache stop
```

Also make sure that it will not be started again after the next reboot by changing the variable `START_HTTPD` in `/etc/rc.config` to “no”. Make sure you run:

```
SuSEconfig
```

to commit the change.

You now need to install the packages with the following commands (assuming the packages reside in the current directory):

```
rpm -Uvh IBM_HTTP_Server-1.3.12-0.i386.rpm  
rpm -Uvh IBM_Admin_Server-1.3.12-0.i386.rpm  
rpm -Uvh IBM_Admin_Server_Forms-1.3.12-0.i386.rpm
```

The installation of the HTTP Server package will also attempt to start the server automatically. If this did not start, you might still have another HTTP Server running. Stop this one first, and try to restart the IBM HTTP Server with the following command:

```
/sbin/init.d/ibmhttpd start
```

If no errors are present on the command line or in the `/opt/IBMHTTPServer/logs/error_log` file, open the new HTTP Server's home page with your browser. You should see the following page:

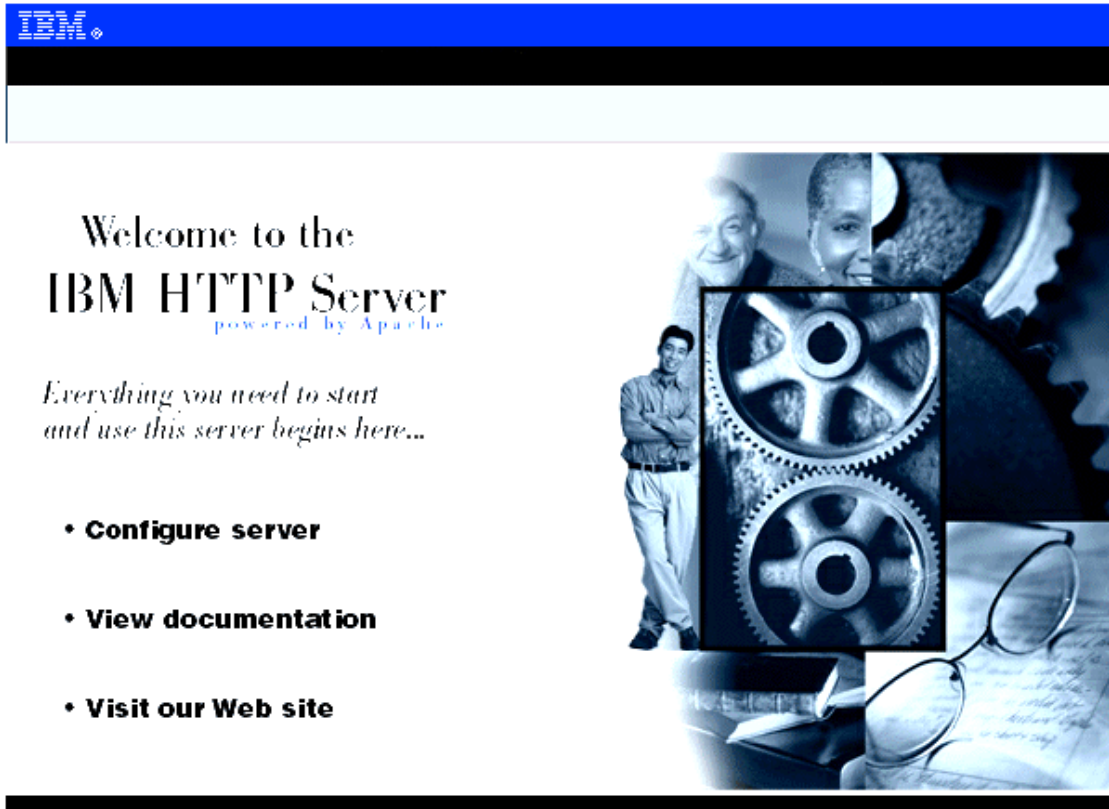


Figure 180. IBM HTTP Server startup page

If you still see the old Web server's startup page (see Figure 179), press Shift+Reload on the Netscape browser to force a reload of this page.

The basic installation of the IBM HTTP Server is now finished. In the default setup, it serves HTML pages from the directory `/opt/IBMHTTPD/htdocs` and CGI scripts from `/opt/IBMHTTPD/cgi-bin`. The log files reside in `/opt/IBMHTTPD/logs`.

7.3.1 Activating IBM HTTPD on system bootup

By default, the IBM HTTP Server has to be started manually after a system reboot. If you want to start it automatically, you have to add the startup script to the bootup procedure. Chapter 17, "The SuSE Linux boot concept" in the SuSE Linux 7.0 manual and the manual page `init.d(7)` give you a detailed description of these mechanisms.

If you want this server to be started on bootup, you have to create the correct symbolic links in the directory `/sbin/init.d/rc2.d` (if you start the system in runlevel 2, the default runlevel), or `/sbin/init.d/rc3.d` (If you use the graphical login, runlevel 3). You can do this manually with the following commands:

```
cd /sbin/init.d/rc2.d
ln -s ../ibmhttpd ./S67ibmhttpd
ln -s ../ibmhttpd ./K01ibmhttpd
```

This will start the IBM HTTP Server in runlevel 2 and makes sure that it will be properly shut down when switching into another runlevel (for example shutdown). Repeat the last two steps above in directory `/sbin/init.d/rc3.d` for runlevel 3 if necessary.

SuSE Linux also ships with a runlevel configuration tool, called `rctab`, which can be used to configure the services to start in this runlevel. To add the script `ibmhttpd` to this runlevel, run `rctab` with the following command line:

```
rctab -e -2
```

This will open an editor (`vi` by default, depending on the environment variable `$EDITOR`) that shows the sequence in which services will be started in this runlevel. Just move to the last entry in the list and add "ibmhttpd" at the first free slot (marked with a "-"). After saving this file, `rctab` will create the necessary symbolic links.

7.3.2 Setting up the administration server

You have to perform some preliminary steps before you can start using the administration server to be able to modify the configuration files of your IBM HTTP Server remotely.

The administration server tasks allow the administration server read/write/execute access to the necessary configuration files and one executable file. The administration server should obtain read/write access through a unique user ID and group, which must be created. The User and Group directives of the administration server's configuration file should be changed to the unique user ID and group. The administration server's configuration file's "group access permissions" should be changed to allow read/write "group access". In addition there is a utility program that should have "Group execute permissions" and "Set User ID Root permissions". This executable must run as root in order to request restarts for the IBM HTTP Server and the administration server.

To properly set up these prerequisites, these tasks can be performed by executing the script `/opt/IBMHTTPserver/bin/setupadm`. After the invocation,

it will ask you a few questions and will give detailed information about each step it is performing. Enter the keywords marked in boldface in the following screens:

```
bash-2.04# ./setupadm

*****
Please supply a User ID to run the Administration Server
We will create the USERID using System Administration tools
*****
[no default] -> wwwrun

*****
Please supply a GROUP NAME to run the Administration Server
We will create the Group using System Administration tools
*****
[no default] -> nogroup

*****
Please supply the Directory containing the files for
which a change in the permissions is necessary.
*****
[default: /opt/IBMHTTPServer/conf] -> [Enter]

These are the file(s) and directory for which we will be changing
Group permissions:

-rw-r--r--  1 root   root      4359 Jun  8 20:46 admin.conf
-rw-r--r--  1 root   root      4359 Jun  8 20:46 admin.conf.default
-rw-r--r--  1 root   root      7453 Jun  8 20:46 admin.msg
-rw-r--r--  1 root   root         1 Jun  8 20:45 admin.passwd
-rw-r--r--  1 root   root     31145 Nov  2 05:45 httpd.conf
-rw-r--r--  1 root   root     31145 Nov  2 05:45 httpd.conf.default
-rw-r--r--  1 root   root     48616 Nov  2 05:45 httpd.conf.sample
-rw-r--r--  1 root   root     12441 Jun  8 20:40 magic
```

Figure 181. Setupadm script

To summarize the above steps: the administration server will be running under the user name “wwwrun” and the group “nobody.”

The administration server is basically just another Web server, running in parallel with the main IBM HTTP Server(s). Therefore it has to be started separately and listens on another TCP port (8008 by default). By default, it has to be started manually. If you also want to start it on system bootup, you have to integrate the start script into the bootup procedure. Copy the file /opt/IBMHTTPServer/bin/adminctl to the directory /sbin/init.d and follow the steps described in 7.3.1, “Activating IBM HTTPD on system bootup” on page 203, using adminctl as the init script name instead of ibmhttpd this time.

The administration server is protected with a user name and password. You can create an entry in the password file /opt/IBMHTTPServer/bin/conf/admin.passwd by issuing the following command from inside the directory /opt/IBMHTTPServer/bin:

```
./htpasswd -m ../conf/admin.passwd <user name>
```

Enter the password for the required user name twice. It is possible to have more than one user name in this password file, if you need to differentiate between multiple administrators.

Now you can start the administration server by running the following command:

```
/opt/IBMHTTPServer/bin/adminctl start
```

After clicking **Configure Server**, shown in Figure 180 on page 203, you need to enter the user name and password you defined for the administration server user. If entered correctly, you will see the welcome page of the administration server:

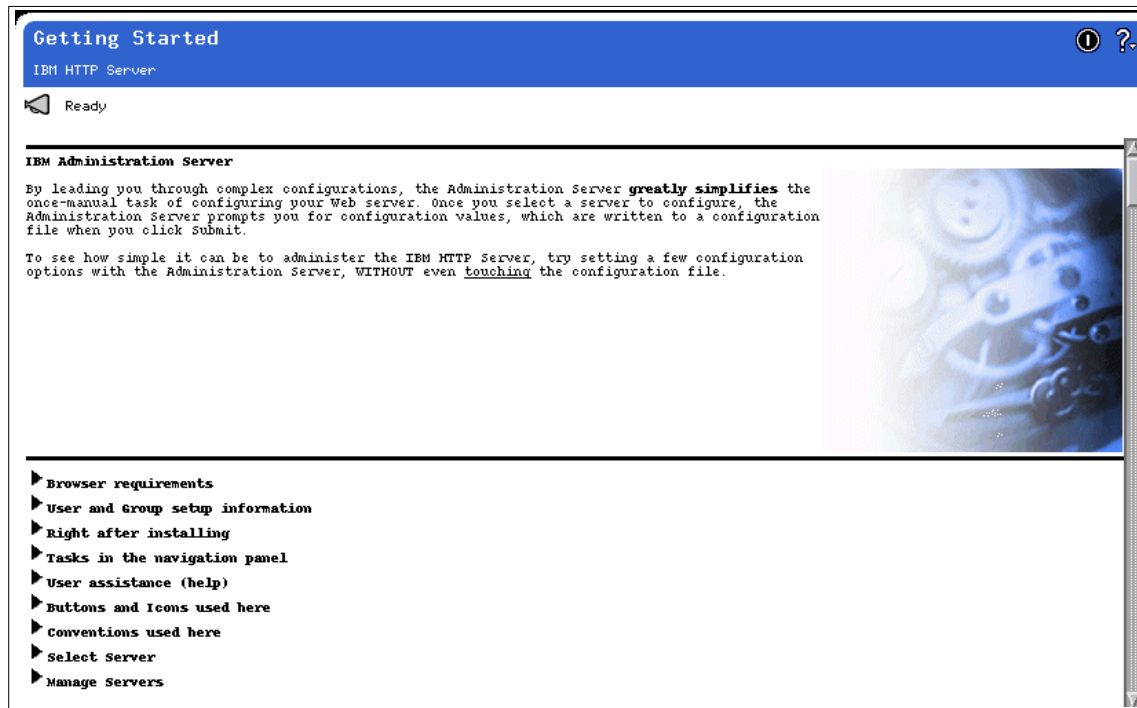


Figure 183. Administration server startup window

You are now ready to start adjusting the configuration of your main Web server according to your needs. Please see the online documentation for help with the different configuration options.

7.4 General performance tips

Configuring Apache for maximum performance is dependent on many parameters. Apache is very flexible and gaining the best performance may require some research. A very informative document about Apache performance tuning can be found on the Apache Web site:

<http://www.apache.org/docs/misc/perf-tuning.html>

In short, experiment with the following options:

- Set the FollowSymLinks option unless you really don't want it.
- Set AllowOverride to None unless you really need it.
- Explicitly list all DirectoryIndex file options from most to least commonly used.

- Tune KeepAliveTimeout starting with 3 ranging to 30 per content and connection types.
- Apache (and the IBM HTTP Server as well) use multiple processes to handle individual requests. Tune StartServers starting with 64 increasing in steps of 32 until performance drops off. Tune MaxClients starting with the value of StartServers. **Note:** Scaling performance can fall off dramatically if Max Clients is too large!
- For SMP systems listening on a single socket, try recompiling after defining SINGLE_LISTEN_UNSERIALIZED_ACCEPT.

A helpful utility to benchmark your Apache server is `ab`. In its simplest form, you can call it like this:

```
ab http://www.your-server.com/index.html
```

The following are `ab` options:

```
Usage: ab [options] [http://]hostname[:port]/path
Options are:
-n requests      Number of requests to perform
-c concurrency   Number of multiple requests to make
-t timelimit     Seconds to max. wait for responses
-p postfile      File containing data to POST
-T content-type  Content-type header for POSTing
-v verbosity     How much troubleshooting info to print
-w              Print out results in HTML tables
-i              Use HEAD instead of GET
-x attributes    String to insert as table attributes
-y attributes    String to insert as tr attributes
-z attributes    String to insert as td or th attributes
-C attribute     Add cookie, eg. 'Apache=1234'. (repeatable)
-H attribute     Add Arbitrary header line, eg. 'Accept-Encoding: zop'
                 Inserted after all normal header lines. (repeatable)
-A attribute     Add Basic WWW Authentication, the attributes
                 are a colon separated username and password.
-p attribute     Add Basic Proxy Authentication, the attributes
                 are a colon separated username and password.
-V              Print version number and exit
-k              Use HTTP KeepAlive feature
-h              Display usage information (this message)
```

Figure 184. `ab` options

Chapter 8. Packet filtering with IP Chains

Whenever you connect your computer to today's Internet world you are exposed to intruders from the outside. There are thousands of hackers just waiting to get into your computer to do damage or maybe to steal information. Therefore you need protection against them!

8.1 What is packet filtering?

As you can tell from the name, packet filtering is a kind of a filter, filtering the data coming to your computer. Packet filtering is one method commonly used in firewall implementations. With packet filtering you can implement a firewall that will protect your computer from the outside world.

Because everybody wants to communicate, sooner or later you need to connect your private network to the Internet. At that point it is time to think about security. You can also use a firewall on a single computer, which is for example connected to the Internet through a dial-up line. When you install a firewall to protect your internal network, every computer that wants to talk to a computer on the internal network must ask the firewall for permission. If the permission is not granted, access is denied.

8.2 What can you do with Linux packet filtering?

With Linux packet filtering you can do many things. Let us describe a few of them here:

- You can protect your internal network connected to the Internet from outside intruders.
- You can perform Network Address Translation (NAT), which allows internally connected computers without a registered Internet address to reach Internet resources.
- You can filter the information going in or out of your internal network or just one computer.
- You can use your Linux server as a gateway between two different types of network, for example connecting token-ring and Ethernet worlds. This can be a cheap solution in comparison to buying an expensive router.
- You can share your dial-up Internet connection with others.

8.3 What do you need to run packet filtering?

To set up a packet filter server with IP Chains, your Linux installation needs to meet requirements:

1. You need kernel Version 2.2.x or higher. It is recommended that you use the latest available stable version. The kernel has to be compiled with appropriate modules for IP Forwarding, IP Masquerading, and IP Firewalling. We recommend that you compile all your networking options and available modules. If you want to use your Linux server as a router, choose **IP - optimize as router not host**. This will also increase the routing performance.
2. Loadable kernel modules Version 2.1.121 or newer
3. IP Chains 1.3.8 or newer

The default installation of SuSE meets all these requirements except that the kernel is not optimized to be used as a router. So if you want to increase the performance of the routing process, you should recompile the kernel and choose the **IP - optimize as router not host** option.

8.4 Network configuration for a packet filtering implementation

In this section we will describe our lab network setup for implementing a packet filtering solution.




Figure 185. Lab network setup for firewall solution

Figure 185 shows our network setup:

- An x230 @server with three Network Interface Cards (NIC) is acting as a gateway. The NICs have the following settings:
 - Eth0 - 192.168.1.1
 - Eth1 - 172.168.1.1
 - Tr0 - 9.24.104.28
- An x340 @server with one NIC and the following settings:
 - Eth0 - 192.168.1.2, default gateway 192.168.1.1
- An x220 @server with one NIC and the following settings:

- Eth0 - 172.168.1.2, default gateway 172.168.1.1
- An NetVista all in one with one NIC and the following settings:
 - Eth0 - 192.168.1.3, default gateway 192.168.1.1

You can see we have two separate networks, 192.168.1.0 and 172.168.1.0. These networks are connected to a Linux server that is acting as a gateway (router). You see that our gateway is connected to the Internet with a registered IP address. We enabled IP Forwarding on the server that was acting as a gateway.

8.5 How to permanently enable IP Forwarding

IP Forwarding is not enabled by default. To enable it, edit the `/etc/rc.config` file and make sure `IP_FORWARD` is set to `yes`. Run `SuSEConfig` to commit the changes and restart the network by typing:

```
rcnetwork restart
```

Now your server is ready to act as a router. You can try this by pinging to the `tr0` interface `9.24.104.28` from the machine on `172.168.1.0` network. If the ping is successful your router is working correctly. You will see a window similar to Figure 186.

```
[root@x220 named]# ping 9.24.104.28
PING 9.24.104.28 (9.24.104.28): 56 data bytes
64 bytes from 9.24.104.28: icmp_seq=0 ttl=255 time=2.5 ms
64 bytes from 9.24.104.28: icmp_seq=1 ttl=255 time=1.1 ms
64 bytes from 9.24.104.28: icmp_seq=2 ttl=255 time=1.0 ms
64 bytes from 9.24.104.28: icmp_seq=3 ttl=255 time=1.0 ms
64 bytes from 9.24.104.28: icmp_seq=4 ttl=255 time=1.2 ms

--- 9.24.104.28 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.0/1.3/2.5 ms
```

Figure 186. PING after enabling IP Forwarding

8.6 Your first IP Chains rule

Now when your router is working, let us make use of it. It does not make sense to have a router without deploying it. We would like to access the external network `9.0.0.0` from the internal network `172.168.1.0`. We can do this by using the IP Masquerading function of IP Chains. Follow these steps on the gateway server to set up the File Transport Protocol (FTP) access from internal network `172.168.1.0` to external network `9.0.0.0`:

1. Create module dependency information for all modules by executing the command:

```
depmod -a
```

2. Load the module for proper FTP masquerading:

```
/sbin/modprobe ip_masq_ftp
```

If you want to use another protocol, such as Real Audio and Internet Relay Chat (IRC), you can load the modules for them also.

3. Set up the timeout for IP Masquerading:

```
/sbin/ipchains -M -S 8000 20 200
```

The parameters have the following meaning:

8000 - timeout value for TCP sessions in seconds

20 - timeout value for TCP sessions after a FIN packet in seconds

200 - timeout value for UDP packets in seconds

You can adjust these settings to meet your needs.

4. Change built-in policy for forwarding by disabling it for all IP addresses:

```
/sbin/ipchains -P forward DENY
```

5. Add the policy for enabling the forwarding with masquerading for your internal networks:

```
/sbin/ipchains -A forward -s 192.168.1.0/24 -j MASQ
```

```
/sbin/ipchains -A forward -s 172.168.1.0/24 -j MASQ
```

You are ready to try your setup. From the computer on the network 172.168.1.0, execute the command:

```
/usr/bin/ftp server
```

Where `ftp server` is the FTP server on the external network (in our example 9.0.0.0). You will see a window similar to Figure 187.

```
[root@x220 named]# ftp 9.24.106.73
Connected to 9.24.106.73.
220 TPIV02 IBM TCP/IP for OS/2 - FTP Server ver 11:45:06 on Apr 17 2000 ready.
Name (9.24.106.73:root): ivo
331 Password required for ivo.
Password:
230 User ivo logged in.
Remote system type is OS/2.
ftp> █
```

Figure 187. FTP after IP Masquerading setup

You have just enabled access from internal networks to an external network.

8.7 How packets travel through a gateway

In this section we will explain how IP Chains work. You can see the path of a packet coming into your server in Figure 188.

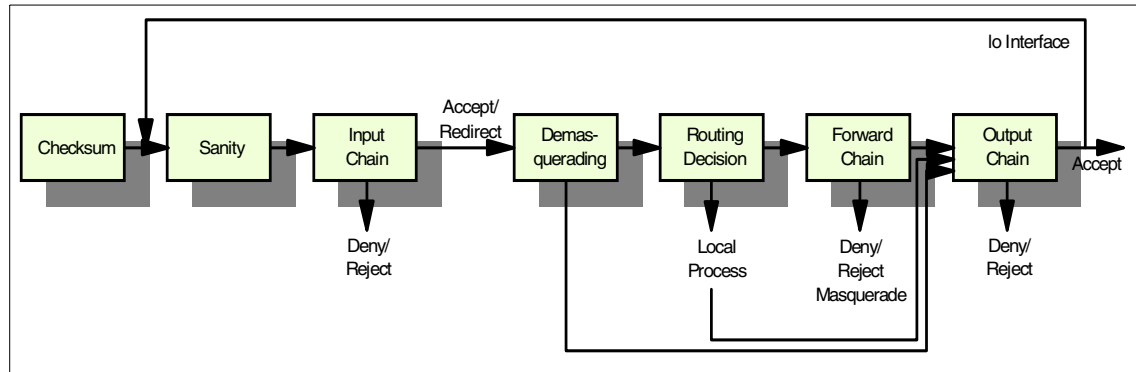


Figure 188. How the packet travels

Here are short descriptions of each stage:

- Checksum - this is to test if the packet is corrupted or not.
- Sanity - Malformed packets are denied here.
- Input chain - This is the first real packet checking point. Packets can be rejected, denied or accepted.
- Demasquerade - If the packet is a reply to a previously masqueraded packet, it is demasqueraded and goes directly from here to the output chain.
- Routing decision - Routing code decides if this packet is for a local process or should be forwarded to a remote machine.
- Local process - a process running on the server can receive packets after a routing decision step, and can then send the packets, which go through a routing decision step and then to the output chain.
- lo interface - if packets from a local process are destined for another local process, they will go through the output chain with interface set to "lo", and will return to the input chain with interface "lo". The "lo" interface is usually called the loopback interface.

- Local - if the packet is not created by the local process, then the forward chain is checked.
- Forward chain - this is the checkpoint for all packets passing through this server to another.
- Output chain - this a checkpoint for all packets just before they are sent out.

As you can see from Figure 188, you have three places where you can check the packets in your server:

- a. Input chain
- b. Forward chain
- c. Output chain

With the `/sbin/ipchains` command you can set up your rules for packet checking.

Note

By default, all checking policies are set to Accept. This means that all packets can come in, go through or go out from your server without any restrictions.

You can see the current checking policies by executing:

```
/sbin/ipchains -L
```

You will see a window similar to Figure 189.

```
[root@client /root]# ipchains -L
Chain input (policy ACCEPT):
Chain forward (policy ACCEPT):
Chain output (policy ACCEPT):
[root@client /root]# █
```

Figure 189. Listing the default IP Chains policies

8.8 Using IP Chains

With the `/sbin/ipchains` command, you can create, change or delete your own policies for checking packets or you can modify built-in policies. You cannot delete the built-in chains, but you can append your rules to the existing chains or even create your own chains.

To manage whole chains you can use the parameters described in Table 17.

Table 17. Parameters for managing whole chains

Parameter	Description
-N	Create a new chain
-X	Delete an empty chain
-P	Change policy for a built-in chain
-L	List the rules in a chain
-F	Flush the rules out of a chain
-Z	Zero the packets and byte counters on all rules in a chain

For manipulating rules inside the chain you can use the parameters explained in Table 18.

Table 18. Parameters for managing rules in the chain

Parameter	Description
-A	Append new rule to a chain
-I	Insert a new rule in a chain at some position
-R	Replace a rule at some position in a chain
-D	Delete a rule at some position in a chain

And there are more operations for managing masquerading. They are described in Table 19.

Table 19. Parameters for managing masquerading

Parameter	Description
-M -L	List the currently masqueraded connections
-M -S	Set masquerading timeout values

8.8.1 How to create a rule

The most common syntax for creating a new rule is:

```
/sbin/ipchains -A input -s source -p protocol -j action
```


The parameters are described in Table 20.

Table 20. IP Chains parameters

Parameter	Description
-A	Append a new rule to the chain
source	IP address or host name of the source
protocol	Type of the protocol to which one a rule is applied
action	What will happen with the packet: 1) ACCEPT - packet will be accepted 2) REJECT - packet will be rejected 3) DENY - packet is dropped since it was not received 4) MASQ - packet will be masqueraded 5) REDIRECT - packet is redirected to local port 6) RETURN - fail off the chain immediately

Note

Redirecting packets to a local port using the REDIRECT action makes sense only in combination with masquerading for a transparent proxy server.

For example, if you want to create a rule for denying the ICMP protocol packets, which are used when you execute the ping command, for a specific IP address you will do this by executing the command:

```
/sbin/ipchains -A input -s IP_address -p icmp -j DENY
```

If you omit the protocol definition, all the packets will be denied. So for example if you want to block the access to your machine from the network 172.168.1.0 with subnet mask 255.255.255.0 you can do this by executing the command:

```
/sbin/ipchains -A input -s 172.168.1.0/255.255.255.0 -j DENY
```

or with:

```
/sbin/ipchains -A input -s 172.168.1.0/24 -j DENY
```

As you can see, the subnet mask can be specified with the number of used bits for that mask.

The command for not allowing any traffic from your server to the network 172.168.1.0 with subnet mask 255.255.255.0 will look like this:

```
/sbin/ipchains -A output -d 172.168.1.0/24 -j DENY
```

Here we used the “-d” parameter for specifying the destination address.

8.8.1.1 Using the inversion flag

With some of the parameters, you can use the inversion option “!”. This means that the rule will be applied to everything else except to the parameters specified after “!”. For example, if you want to deny packets that come from all IP addresses except from network 192.168.1.0 with subnet mask 255.255.255.0 you can do this by executing the command:

```
/sbin/ipchains -A input -s ! 192.168.1.0/24 -j DENY
```

Note

The rules you made are not permanent, so next time you restart the server they will be lost.

8.8.2 Making the rules permanent

For making the rules permanent you have two scripts available that can make your life easier. To save all the rules you created, you can execute the command:

```
/sbin/ipchains-save > filename
```

If you execute this command without a file name, the rules will be sent to the standard output.

You can then restore the saved rules by executing the command:

```
cat filename | /sbin/ipchains-restore
```

So if you want your saved rules to be enabled whenever you start your system, add the following line to the `/etc/rc.d/rc.local` file:

```
cat filename | /sbin/ipchains-restore
```

8.9 Sources of additional information

You can find more information on the official Linux IP Firewall Chains page at:

```
http://www.rustcorp.com/linux/ipchains
```

And there are always good how-to documents on the Linux Documentation Project home page:

```
http://www.linuxdoc.org/
```

Chapter 9. DHCP - Dynamic Host Configuration Protocol

With the ever-decreasing number of IP addresses available, along with the headache of maintaining static IPs, DHCP has become a necessity in most TCP/IP computing environments.

9.1 What is DHCP?

DHCP stands for Dynamic Host Configuration Protocol. When using TCP/IP, a computer system needs a unique IP address to communicate with other computer systems. Without DHCP, the IP address must be entered manually at each computer system. DHCP lets network administrators distribute IP addresses from a central location without having to actively manage each individual address.

With DHCP, IP addresses are distributed through pools usually broken up by subnet. Leases are given out for a specific time period for each address. The process of managing leases is all done by the DHCP server. Once a lease has expired the DHCP server will try to contact the client or the client will contact the server to renew the lease. If the server cannot contact the client, the IP address is returned to the pool and available for the next client in need of an address.

9.2 Why should I use DHCP?

In the past, for every device on a network you had to have a static IP address. With the increasing number of computers accessing the Internet, the pool of available addresses is quickly diminishing. Network administrators can significantly reduce the number of IP addresses they need by using DHCP.

Even with smaller networks, keeping track of individual IP addresses can be maintenance intensive. With DHCP, the server does all of the maintenance, mapping IP addresses to MAC addresses and tracking lease times. Administrators can adjust lease times, expand or reduce pools, and change gateways or DNS addresses, all from a central location.

9.3 Implementation on Linux

In this section we will discuss how to implement a DHCP server on Linux.

The DHCP server can be found in the network (n) series. Use YaST to install it.

Using your editor of choice, create the `/etc/dhcpd.conf` file.

The following sample `dhcpd.conf` file is rather simple. We designate a default lease time of 600 seconds (10 minutes) but we will let clients request up to a 7200-second (2-hour) lease time. We include a recommended subnet mask of 255.255.255.0 and a broadcast address of 192.168.1.255. Other options we specify include a default gateway (router), two name servers, and the domain.

For our subnet specifics we are using the private 192.168.1.0 class C subnet. For our DHCP pool we will be giving out addresses numbered from 15 to 100 for a total of 85 addresses. The rest can be used by static devices.

```
#!/etc/dhcpd.conf
default-lease-time 600;
max-lease-time 7200;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option routers 192.168.1.254;
option domain-name-servers 192.168.1.1, 192.168.1.2;
option domain-name "ibm.com";

subnet 192.168.1.0 netmask 255.255.255.0 {
    range 192.168.1.15 192.168.1.100;
}
```

You are not limited to a single subnet. You are allowed to have shared network-specific parameters, multiple subnet-specific parameters, group parameters, and host-specific parameters.

You can define multiple ranges, assign specific IP addresses based on the hardware address of the client, and specify a WINS server if needed.

More information is available from the `dhcpd.conf` man page.

To start the DHCP daemon, type:

```
rcdhcpd start
```

Chapter 10. Postfix

Sendmail is considered the de-facto mail system for UNIX systems. It includes all the features needed in a mail system and has been in circulation for many years, and as such provides excellent stability.

The only problem with Sendmail is the configuration of the service. Most Linux services use a common configuration file format, commonly known as rc files. rc files use familiar *variable=value pairs* that are usually self explanatory.

Sendmail, on the other hand, uses cryptic variable/value pairs. Having a system that is difficult to configure and administer can lead to the administrator overlooking or wrongly setting configuration options. It also affects the time it takes to deliver the system, which can be a big decider in choosing a system configuration.

Postfix was written by Wietse Venema while he was a resident at an IBM research facility. In Wietse's own words:

“Postfix is my attempt to provide an alternative to the widely used Sendmail program. Postfix attempts to be fast, easy to administer, and hopefully secure, while at the same time being sendmail-compatible enough to not upset your users.”

10.1 Installing Postfix

The installation of Postfix isn't as straightforward as other applications, since you will need to remove the current mail system installed. By default Sendmail is installed in a SuSE Linux system. Both Sendmail and Postfix can be found in the **n** (networking) section of the package selection dialog in YaST.

You will need to remove Sendmail before installing Postfix, since both systems will “battle” to run and will produce unpredictable results and will almost definitely prohibit one of the servers from running. If you do not remove Sendmail, or do not set Sendmail to be removed, you will receive a warning from YaST that it is not advisable for both of the services to be running on the same system.

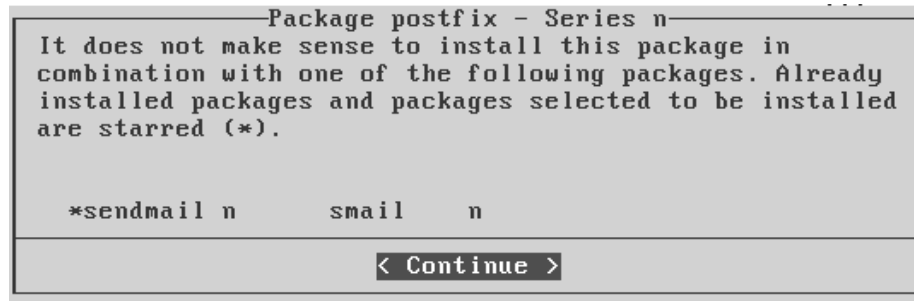


Figure 190. Postfix and Sendmail on the same system

10.1.1 Postfix configuration files

Once you have installed Postfix you can start to configure it. Postfix's configuration files are based in the `/etc/postfix` directory. The most important configuration file is `/etc/postfix/main.cf`. Other configuration files are:

access - This configuration file dictates what hosts may access resources through Postfix.

canonical - This defines address remapping rules, both local and remote.

main.cf - This file is the global configuration file for Postfix.

master.cf - Defines the master daemon configuration that controls Postfix daemons for incoming and outgoing mail.

relocated - Defines members of the organization that are no longer using the mail system. Postfix will notify senders that the recipient is no longer at the address and can also tell the sender of a new e-mail address where the recipient can be contacted.

More information on these files and their configuration options can be found by typing:

```
man filename
```

Where `filename` is the name of the Postfix configuration file about which you want more information.

10.2 Example Postfix installation

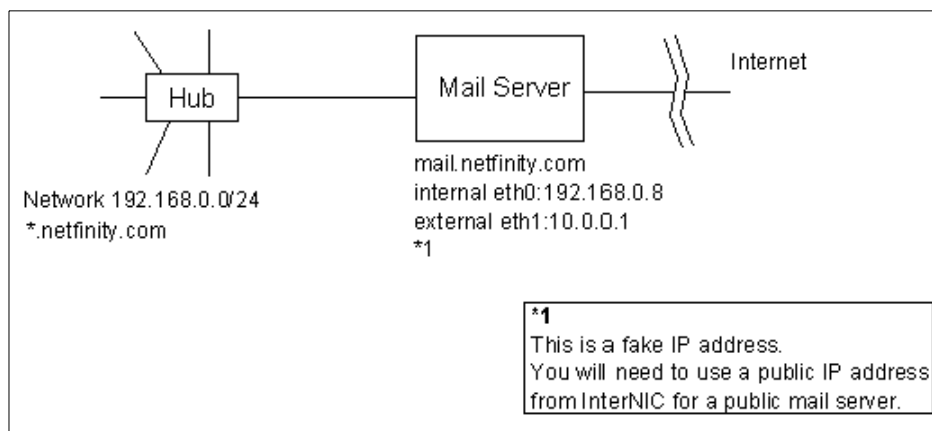


Figure 191. Mail server test configuration

The mail system will be configured to provide local mail services for the internal network (*.netfinity.com) and provide mail services for incoming Internet mail. The mail server interface that will be connected to the Internet will need a public IP address, and will also need to have a registered MX record in DNS. The MX record can either be handled by your DNS server if you have one, or you can use your ISP to handle the DNS records for the mail server.

In our situation, once the MX records have been updated all mail for justin@netfinity.com will be directed to mail.netfinity.com automatically by other mail servers looking up the MX handle in the DNS record for the netfinity.com domain.

10.2.1 Telling Postfix about the server

The first thing we need to do is tell Postfix what domain it is working on. Postfix works from the FQDN of the mail server, in our case mail.netfinity.com. The FQDN will be automatically adjusted in different situations of the Postfix configuration, the most common being to truncate the host part of the FQDN to get the domain name of the server.

The main Postfix configuration file, /etc/postfix/main.cf, is a placeholder for the configuration options we choose to edit. Use your favorite text editor to edit the file as detailed in the rest of the chapter.

```

# INTERNET HOST AND DOMAIN NAMES
#
# The myhostname parameter specifies the internet hostname of this
# mail system. The default is to use the fully-qualified domain name
# from gethostname(). $myhostname is used as a default value for many
# other configuration parameters.
#
myhostname = mail.netfinity.com

```

Figure 192. Internet host and domain names section of the main.cf file

This is usually automatically set for you when SuSE Linux installs postfix.

10.2.2 Telling the rest of the world about us

The mail server needs to know where it is supposed to state where the mail was originated. By default it uses *myhostname*, but this is not the preferred way on the Internet. We need to specifically add this identifier to the configuration as the installation uses the Postfix default setting.

```

$myorigin= $mydomain

```

Figure 193. Myorigin section of the main.cf file

We need to tell Postfix what domains we will handle mail for. By default it is not set. Edit main.cf with the following in **bold** added.

```

# The local machine is always the final destination for mail addressed
# to user@[the.net.work.address] of an interface that the mail system
# receives mail on (see the inet_interfaces parameter).
#
# Specify a list of host or domain names, /file/name or type:table
# patterns, separated by commas and/or whitespace. A /file/name
# pattern is replaced by its contents; a type:table is matched when
# a name matches a lookup key. Continue long lines by starting the
# next line with whitespace.
#
# DO NOT LIST VIRTUAL DOMAINS HERE. LIST THEM IN THE VIRTUAL FILE
# INSTEAD. BE SURE TO READ THE ENTIRE VIRTUAL MANUAL PAGE.
#
mydestination = $myhostname, localhost.$mydomain
mydestination = $myhostname, localhost.$mydomain $mydomain
mydestination = $myhostname, localhost.$mydomain, $mydomain,
#      mail.$mydomain, www.$mydomain, ftp.$mydomain
mydestination = $myhostname, localhost.$mydomain, $mydomain

```

Figure 194. Mydestination section of the main.cf file

We have added mydomain to the list of hosts that this machine considers itself the final destination for. We have told Postfix that it should consider all mail that passes through the system with the destination as @netfinity.com as local. In this case the mail will be handed over to the local mail agent that will take care of delivering the mail to the local system.

We need to “hide” the hosts behind the mail server from the rest of the Internet. This is not the same masquerading that is provided by firewalls or masquerade servers. This type of masquerading allows the address justin@client121.netfinity.com to be translated to justin@netfinity.com on the way through the mail server.

To enable this feature you will need to add the following statement:

```
masquerade_domains = $mydomain
```

Figure 195. Masquerade_domains section of the main.cf file

to the main.cf file. It is not enabled by default. This will truncate the sender address coming from the internal network to netfinity.com.

10.2.3 Transport and delivery options

```
# DEFAULT TRANSPORT
#
# The default_transport parameter specifies the default message
# delivery transport to use when no transport is explicitly given in
# the optional transport(5) table.
#
default_transport = smtp
#default_transport = uucp
```

Figure 196. Default_transport section of the main.cf file

In the old days UUCP was used to transfer mail between UNIX systems. It stands for UnixtoUnixCopy and is still used in some situations today. As we are using the Internet as our network handler, we should specify Simple Mail Transfer Protocol (SMTP) as our mail transport protocol.

```

# The mailbox_command specifies the optional external command to use
# instead of mailbox delivery. The command is run with proper HOME,
# SHELL and LOGNAME settings.
#
# Avoid shell meta characters because they will force Postfix to run
# an expensive shell process. Procmail alone is expensive enough.
#
# IF YOU USE THIS TO DELIVER MAIL SYSTEM-WIDE, YOU MUST SET UP AN
# ALIAS THAT FORWARDS MAIL FOR ROOT TO A REAL USER.
#
mailbox_command = /usr/bin/procmail
# mailbox_command = /usr/bin/procmail -a "$EXTENSION"

```

Figure 197. Mailbox_command section of the main.cf file

Delivery of mail on the local system is usually uneventful in nature. You can use a rule-based mechanism to sort mail per user. The recipient is also able to use this mechanism in deciding, for example, that mail from a mailing list should be saved in a separate mail file.

The program in question is procmail. It allows the user to store a file called .procmailrc (not the full stop) that contains rules for incoming mail. The command set for procmail commands is too big to be described here. Take a look at the man procmailrc file for a description of what can be done with procmail.

10.2.4 Security goes a long way

```

# The mynetworks parameter specifies the list of networks that are
# local to this machine. The list is used by the anti-UCE software
# to distinguish local clients from strangers. See permit_mynetworks
# and smtpd_recipient_restrictions in the file sample-smtpd.cf file.
#
# The default is a list of all networks attached to the machine: a
# complete class A network (X.0.0.0/8), a complete class B network
# (X.X.0.0/16), and so on. If you want stricter control, specify a
# list of network/mask patterns, where the mask specifies the number
# of bits in the network part of a host address. You can also specify
# the absolute pathname of a pattern file instead of listing the
# patterns here.
#
mynetworks = 192.168.0.0/24, 10.0.0.0/24, 127.0.0.0/8
#mynetworks = $config_directory/mynetworks

```

Figure 198. Mynetworks section of the main.cf file

The mynetworks parameter among other things is used in deciding what hosts are allowed to relay mail through the mail server. Mail relaying is a

feature that all mail servers have, but is the biggest pitfall for system administrators, because relaying is used by mail spammers to send bulk mail to the Internet or to hide the identity of malicious users.

If you allow relaying from all hosts you may end up getting blacklisted by recipients ISPs, if someone sends “spam” from your host to its customers. This will result in mail to that ISP’s customers (it can be a considerable amount of users if it is AOL, Freeserve or any of the other big ISPs) being blocked from your mail server. You have been warned; it is a very common thing to see in your logs that someone has tried to use your mail server for relaying. Do not think it will not happen to you. It will.

Extended Turn (ETRN) commands are not needed in our configuration. By default Postfix allows all hosts to issue ETRN commands to the server. This is unnecessary and can present security holes in the system. We need to restrict access to ETRN commands in our system.

```
smtpd_etrn_restrictions= $mynetworks, reject
```

Figure 199. *Smtpd_etrn_restrictions* section of the *main.cf* file

10.2.5 Reporting Postfix errors

A mail server is one of the busiest services employed on a system, and as such should be monitored by the administrator for erratic behavior.

The log file for the mail server is located in `/var/log/mail`. If you need to troubleshoot the server, this is the first place to look for errors.

The administrator should also be kept up to date about errors during the running of Postfix. We can set up the server to send an e-mail to the *post master* user when an error or a trigger occurs.

```
notify_classes = bounce, 2bounce, delay, policy, protocol, resource, software
```

Figure 200. *Notify_classes* section of the *main.cf* file

The notice triggers are:

bounce - Notify the administrator that an e-mail has bounced.

2bounce - Notify the administrator that an e-mail has double bounced.

delay - Notify the administrator that an e-mail has been delayed during delivery.

policy - Notify the administrator of rejected messages as a result of policies enforced by Postfix and the administrator.

protocol - Notify the administrator of an error related to the mail protocol, for example malformed mail requests.

resource - Notify the administrator that a message could not be delivered because of resource errors, for example low disk space.

software - Notify the administrator about software errors, for example an error in the configuration of Postfix.

Chapter 11. NFS - Network File System

The Network File System (NFS), developed by Sun Microsystems, allows you to share directories across the network. The directory mounts become transparent to you. You access the mounted directories just as you do any directory or file system on your computer. The mounting process is the same as for any file system or partition that you want to mount on your system. The basic foundation of this is the mount command.

In order to share directories across the network you will need two basic things:

- The system sharing the data must allow you to have access
- The system that is using the data must originate the request and allow the mount to happen

Both concepts will be discussed in this chapter.

11.1 The NFS process

First, you need to verify that the NFS package has been loaded. To use the RPM package. RPM package is short for Red Hat Package Manager and is a common way of installing packages in Linux. You can do this with the command:

```
rpm -q knfsd
```

If the NFS server has been installed you will be shown the package version. If it can not be found, install the knfsd package as detailed in 3.1, “Adding and removing software packages using YaST” on page 111. The knfsd package can be found in the n (networking) series.

NFS makes use of several daemons. Those daemons are:

- | | |
|-------------|--|
| portmap | This is the process that converts Remote Procedure Call (RPC) program numbers into Defense Advanced Research Projects Agency (DARPA) protocol port numbers. When a client wishes to make an RPC call to a given program number (for example, the NFS server), it will first contact portmap on the server machine to determine the port number where RPC packets should be sent. |
| rpc.kmountd | This handles the exporting of NFS file systems. It looks in the /etc/exports file to figure out what to do with mount requests from various hosts. |

`rpc.knfsd` This provides the user level part of the NFS process.

`rpc.rquotad` This handles quotas for access to file systems. The quotas are based on disk usage and can be hard or soft limits.

You can verify that the `rpc.knfsd`, `rpc.kmountd`, and `portmap` daemons are running as shown in Figure 201.

```
# ps ax | grep nfs
323 ?      SW      0:00 [nfsd]
324 ?      SW      0:00 [nfsd]
325 ?      SW      0:00 [nfsd]
326 ?      SW      0:00 [nfsd]
327 ?      SW      0:00 [nfsd]
328 ?      SW      0:00 [nfsd]
329 ?      SW      0:00 [nfsd]
330 ?      SW      0:00 [nfsd]
662 ttypl  S        0:00 grep nfs
#
# ps ax | grep mount
313 ?      SW      0:00 [rpc.kmountd]
673 ttypl  S        0:00 grep mount
#
# rportmap status
OK
```

Figure 201. Verifying the NFS daemons

If the `portmap` daemon is not running, you need to start it up first before you start up the NFS daemons. You can do this with the command:

```
rportmap start
```

Once the `portmap` daemon is running, you can start up the NFS daemons with the command:

```
rcnfsserver start
```

```
bash-2.04# rcnfsserver start
Starting kernel based NFS server
```

Figure 202. Starting up NFS

Note

If the `/etc/exports` file does not exist or is empty, the NFS daemons will not start. Information on setting up the `/etc/exports` file is in 11.2, “Allowing NFS access to data” on page 233.

To stop the NFS server, use the command:

```
rcnfsserver stop
```

You can restart the NFS process with the command:

```
rcnfsserver restart
```

This can also be used to restart the NFS process if you have made changes to the configuration files.

11.2 Allowing NFS access to data

You can give NFS access to a file system by setting it up in the `/etc/exports` file. The file is set up on the exporting server. You can create a sample file entry by opening the `/etc/exports` file. Then you can add an entry like:

```
/usr/local/share myserver.mydomain.com(ro)
```

This says that the directory `/usr/local/share` is only accessible to the server `myserver.mydomain.com`.

Note

When exporting a file system you need to be sure that the exporting server can recognize and access the server that is in the `/etc/exports` file. You can verify this with the command:

```
ping server_name
```

Where `server_name` is the name of the server you are trying to access. Otherwise, the NFS commands may hang.

There are a number of options you can set up in the `/etc/exports` file. Some of them are listed in Table 21.

You need to be sure that the exporting server can recognize the server name.

The various options are explained in the table below.

Table 21. Access options

ro read only	Only permits reading
rw read write	Permits reading and writing. If both ro and rw are specified, rw takes priority.
root_squash client	Anonymous user (nobody) access from client.
no_root_squash client	Access request privileges per the privileges of the client root. Useful for diskless clients.
squash_uids and squash_gids	Specify a list of UIDs or GIDs that should be subject to anonymous mapping. A valid list of IDs looks like this: squash_uids=0-15,20,25-50
all_squash all access	Processes all requests for access as anonymous user.
anonuid=uid	root_squash or all_squash when options are set will assign a group ID to an anonymous user request.
anonuid=gid	root_squash or all_squash when options are set will assign a group ID to an anonymous user request.

A sample `/etc/exports` file is shown in the man pages for `exports` and in Figure 203.

```
# sample /etc/exports file
/          master(rw) trusty(rw,no_root_squash)
/projects  proj*.local.domain(rw)
/usr       *.local.domain(ro) @trusted(rw)
/home/joe  pc001(rw,all_squash,anonuid=150,anongid=100)
/pub       (ro,insecure,all_squash)
/pub/private (noaccess)
```

Figure 203. A sample `/etc/exports` file

The lines in the sample `/etc/exports` file are explained as follows:

```
# sample /etc/exports file
```

This is just a comment. Any line or character string can be converted to a comment and disabled by entering a `#` symbol. Everything from that point to the end of the line is considered to be a comment.

```
/ master(rw) trusty(rw,no_root_squash)
```

This says that the root directory (`/`) is exported to the servers:

master - whose rights are read-write

trusty - whose rights are read-write and the access rights of the client root can be the same as the server's root

```
/projects      proj*.local.domain(rw)
```

The directory /projects is read-write accessible to all servers whose names match the pattern `proj*.local.domain`. This includes `proj.local.domain`, `proj1.local.domain`, `projprojproj.local.domain` and so forth.

- /usr *.local.domain(ro) @trusted(rw)

Any systems whose hostname ends in `.local.domain` is allowed read-only access. The `@trusted` netgroup is allowed read-write access.

- /home/joe pc001(rw,all_squash,anonuid=150,anongid=100)

The directory /home/joe is accessible to pc001 for read-write access; all requests for access are processed as anonymous users. The anonymous UID number is set to 150 and the anonymous group ID is set to 100. This is useful when using a client that is running PCNFS or an equivalent NFS process on the PC. Since the PC IDs do not necessarily map to the UNIX IDs, this allows the proper file attributes to be set.

- /pub (ro,insecure,all_squash)

The directory /pub is accessible as read-only. The option in this entry also allows clients with NFS implementations that don't use a reserved port for NFS and process all requests as an anonymous user.

- /pub/private (noaccess)

The directory /pub/private does not allow any NFS access.

11.3 Accessing data remotely with NFS - the command line view

To mount a remote file system on your local system, the mount point must exist. The mount process does not create the mount point automatically. The process of making the mount point is to use the Linux `mkdir` command. To make the `/usr/local/share` mount point, enter:

```
mkdir /usr/local/share
```

Typically you do not need to worry about file attributes and ownership when making an NFS mount point. The NFS access rights will usually supersede any rights established for the directory.

Once you have created the mount point, you can use the mount command as follows:

```
mount -t nfs nfs_host:share_dir local_mount_dir
```

where:

<code>-t nfs</code>	Says to do the mount as an NFS mount. This is now optional because if you explicitly specify the directory to be mounted as <code>host:directory</code> the <code>mount</code> command knows that it is an NFS mount.
<code>nfs_host</code>	Is the host that is exporting the file system to be shared.
<code>share_dir</code>	Is the actual directory that is to be shared.
<code>local_mount_dir</code>	Is the directory on the local host where the remote directory is going to be mounted. As mentioned earlier, this mount point must exist.

Chapter 12. NIS - Network Information System

In a distributed computing environment, maintenance of password, group, and host files can be a major task. Consistency is possibly the biggest difficulty here. For example, when a user changes his password on one machine, ideally it would be propagated to any other machine he has accounts on. When a network is composed of hundreds or thousands of machines, this convenience becomes a necessity. NIS is one way of addressing some of these problems.

12.1 What is NIS?

The Network Information System (NIS) is a service designed to provide a distributed database system for common configuration files. It was formerly known as Sun Yellow Pages (YP). NIS servers manage copies of the database files. NIS clients request the information from the NIS server instead of using their own configuration files.

NIS is designed after the client/server model. A NIS server contains data files called "maps". These maps are owned by the NIS master and can only be updated by the master. There are NIS slave servers that replicate from the master. When there is a change to a master server's map, this change is then distributed to all the slave servers. Clients are hosts that request information from these maps but are not allowed to modify them locally.

NIS is commonly used in UNIX environments. However, it is also possible to integrate Windows NT clients in a NIS-based environment. NISGINA provides a NIS authenticated interactive logon for Windows NT 4.0 workstations. It supports changing UNIX passwords using a Windows NT dialog and some limited remote registry configuration. You can find it at the author's Web page at:

<http://www.dcs.qmw.ac.uk/~williams/>

12.2 How can I use NIS?

NIS is typically used to centrally manage commonly replicated configuration files. Examples of common configuration files are:

- /etc/hosts
- /etc/passwd
- /etc/group

12.3 Implementation on Linux

To introduce the concepts behind NIS, we will create a map of our password file kept on the NIS master server. This will allow users to log in to NIS clients without having to maintain an account on each system. Centralized administration is a key benefit of using NIS.

A note on security: Before deciding to put NIS in a production environment, please consider the security implications of passing sensitive data across the network. You may wish to take a look at NIS+, which has strong encryption as well as additional maintenance implications.

The `ypclient` packages needs to be installed for a NIS client.

Packages that need to be installed for a NIS server:

- `ypserv`
- `make`

The `nsswitch` file determines the order of lookups performed. In the case of the `passwd` entry, the system will query the local password database, and then query the NIS server for the relevant username entry.

Sample `nsswitch.conf` file:

```
# /etc/nsswitch.conf

passwd: files nis
shadow: files nis
group:   files nis
hosts:   files dns
bootparams: files
ethers:  files
netmasks: files
networks: files
protocols: files
rpc:     files
services: files
#netgroup: nisplus
#publickey: nisplus
automount: files
aliases:  files
```

12.3.1 NIS server

A key configuration file for the NIS master server is the `/etc/ypserv.conf` file. Uncomment the `passwd.byname` line, but leave the `passwd.byuid` commented. The following is a sample `ypserv.conf` we used:

```
#ypserv.conf - In this file you can set certain options for the NIS
server,
#and you can deny or restrict access to certain maps based
#on the originating host.
#See ypserv.conf(5) for a description of the syntax.
dns: no
# The following, when uncommented, will give you shadow like passwords.
# Note that it will not work if you have slave NIS servers in your
# network that do not run the same server as you.
# Host          : Map          : Security   : Passwd_mangle
*               : passwd.byname  : port       : yes
# *             : passwd.byuid   : port       : yes
# Not everybody should see the shadow passwords, not secure, since
# under MSDOS everybody is root and can access ports < 1024 !!!
* : shadow.byname  : port       : yes
* : passwd.adjunct.byname : port       : yes
# If you comment out the next rule, ypserv and rpc.ypxfrd will
# look for YP_SECURE and YP_AUTHDES in the maps. This will make
# the security check a little bit slower, but you only have to
# change the keys on the master server, not the configuration files
# on each NIS server.
# If you have maps with YP_SECURE or YP_AUTHDES, you should create
# a rule for them above, that's much faster.
*               : *           : none
```

The other key configuration file is the `/var/yp/Makefile`. The only map we want to create is the `/etc/passwd` file, so the others can be commented out if you wish; however, the default `Makefile` works just fine. The following is a sample `/var/yp/Makefile`:

```
# Makefile for the NIS databases
# This Makefile should only be run on the NIS master server of a domain.
# All updated maps will be pushed to all NIS slave servers listed in the
# /var/yp/ypservers file. Please make sure that the hostnames of all
# NIS servers in your domain are listed in /var/yp/ypservers.
# This Makefile can be modified to support more NIS maps if desired.
# Set the following variable to "-b" to have NIS servers use the domain
# name resolver for hosts not in the current domain. This is only
# needed,
# if you have SunOS slave YP server, which gets here maps from this
# server. The NYS YP server will ignore the YP_INTERDOMAIN key.
#B=-b
```

```

B=
# If we have only one server, we don't have to push the maps to the
# slave servers (NOPUSH=true). If you have slave servers, change this
# to "NOPUSH=false" and put all hostnames of your slave servers in the
# file
# /var/yp/ypservers.
NOPUSH=true
# We do not put password entries with lower UIDs (the root and system
# entries) in the NIS password database, for security. MINUID is the
# lowest uid that will be included in the password maps.
# MINGID is the lowest gid that will be included in the group maps.
MINUID=500
MINGID=500
# Should we merge the passwd file with the shadow file ?
# MERGE_PASSWD=true|false
MERGE_PASSWD=true
# Should we merge the group file with the shadow file ?
# MERGE_GROUP=true|false
MERGE_GROUP=true
# These are commands which this Makefile needs to properly rebuild the
# NIS databases. Don't change these unless you have a good reason.
AWK = /usr/bin/gawk
MAKE = /usr/bin/gmake
UMASK = umask 066
# These are the source directories for the NIS files; normally
# that is /etc but you may want to move the source for the password
# and group files to (for example) /var/yp/ypfiles. The directory
# for passwd, group and shadow is defined by YPPWDDIR, the rest is
# taken from YPSRCDIR.
YPSRCDIR = /etc
YPPWDDIR = /etc
YPBINDIR = /usr/lib/yp
YPSBINDIR = /usr/sbin
YPPDIR = /var/yp
YPMAPDIR = $(YPPDIR)/$(DOMAIN)
# These are the files from which the NIS databases are built. You may
# edit
# these to taste in the event that you wish to keep your NIS source files
# separate from your NIS server's actual configuration files.
GROUP      = $(YPPWDDIR)/group
PASSWD     = $(YPPWDDIR)/passwd
SHADOW     = $(YPPWDDIR)/shadow
GSHADOW    = $(YPPWDDIR)/gshadow
ADJUNCT    = $(YPPWDDIR)/passwd.adjunct
#ALIASES   = $(YPSRCDIR)/aliases # aliases could be in /etc or
/etc/mail
ALIASES    = /etc/aliases

```

```

ETHERS      = $(YPSRCDIR)/ethers      # ethernet addresses (for rarpd)
BOOTPARAMS  = $(YPSRCDIR)/bootparams # for booting Sun boxes
            (bootparamd)
HOSTS       = $(YPSRCDIR)/hosts
NETWORKS    = $(YPSRCDIR)/networks
PROTOCOLS   = $(YPSRCDIR)/protocols
PUBLICKEYS  = $(YPSRCDIR)/publickey
RPC         = $(YPSRCDIR)/rpc
SERVICES    = $(YPSRCDIR)/services
NETGROUP    = $(YPSRCDIR)/netgroup
NETID       = $(YPSRCDIR)/netid
AMD_HOME    = $(YPSRCDIR)/amd.home
AUTO_MASTER = $(YPSRCDIR)/auto.master
AUTO_HOME   = $(YPSRCDIR)/auto.home
YPSERVERS  = $(YPDIR)/ypservers# List of all NIS servers for a domain
target: Makefile
@test ! -d $(LOCALDOMAIN) && mkdir $(LOCALDOMAIN) ; \
cd $(LOCALDOMAIN) ; \
$(NOPUSH) || $(MAKE) -f ../Makefile ypservers; \
$(MAKE) -f ../Makefile all
# If you don't want some of these maps built, feel free to comment
# them out from this list.
all: passwd group hosts rpc services netid protocols netgrp mail \
#shadow publickey # networks ethers bootparams amd.home \
auto.master auto.home passwd.adjunct
#####
#
# DON'T EDIT ANYTHING BELOW IF YOU DON'T KNOW WHAT YOU ARE DOING !!! #
#####
#

```

Take a look at the `/var/yp/securenets` file, which defines the access rights to your NIS server. By default it is set to give access to everyone. Change it accordingly (see the man `securenets` file).

We need to set our nisdomain name. This domain name should not be confused with DNS domain names! The YP domain name can be any generic name.

Edit the `/etc/rc.config` file, and change the `YP_DOMAINNAME` value to “`nis.com`” for our purposes. As usual, run `SuSEconfig` to commit the changes.

You also have to define which hosts should be allowed to contact the NIS server. In our example, we will allow all hosts from the local Class C network `192.168.158.0/24` to connect to the server.

Open `/etc/hosts.allow` in a text editor and add the following line:

```
ypserv: 127.0.0.0/255.0.0.0 192.168.158.0/255.255.255.0
```

It is imperative that the local host also be allowed to connect to the `ypserv` process via the loopback interface (`127.0.0.1`).

Now add the following line to `/etc/hosts.deny`:

```
ypserv: ALL
```

We are now ready to start the `ypserv` daemon:

```
rcypserv start
```

To test our NIS setup we will use the `rpcinfo` command:

```
rpcinfo -u localhost ypserv
```

You should see:

```
program 100004 version 1 ready and waiting
program 100004 version 2 ready and waiting
```

We will now create our NIS maps:

```
/usr/lib/yp/ypinit -m (to create the DB)
(ypinit -s masterhost to add a slave server.)
```

The localhost is selected as the master server:

```
<ctrl> d
```

Select `y` to confirm and begin building your maps.

To test our NIS master server, we need to set up a client to run `ypbind`. For simplicity we can use the master server to verify our configuration. The same steps should be followed to set up a remote client.

12.3.2 NIS Client

We have to set up the client to access the NIS server and request information for the local system. To do this, load YaST from a command prompt and select **System Administration->Network Configuration->Configure YP Client**.

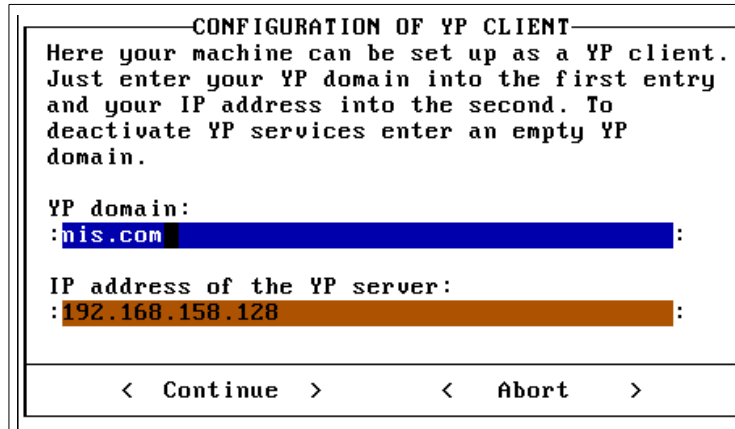


Figure 204. Configure YP client

Enter the YP domain name. This is the same as the YP domain that we entered for the server, as this is the NIS server we will be requesting information from. In our case it will be “nis.com” (no quotes). You will also need to enter the IP address of the YP server. Press **Continue** to allow YaST to commit your settings.

Once YaST has completed SuSEconfig, exit and type:

```
rcypclient start
```

This will start the ypbind daemon and connect to the YP server. You are now able to access the NIS server.

To test our NIS configuration we will use the `ypcat` command:

```
ypcat passwd
```

You should see output similar to:

```

[root@m10A /root]# ypcat passwd
jay:$1$FWGys23z$oUHJSt7bvnc9rP15xvb/w.:500:500:~/home/jay:/bin/bash
ayne:$1$34Gys23zsjdY0Gnc9rP15xvb/w.:501:501:~/home/ayne:/bin/bash
kathy:$1$gJghws23zsjdYJGnc65b/g5.:502:502:~/home/kathy:/bin/bash
osa:$1$FWFvdggJw7bvnc9rP15xgdhd/4.:503:503:~/home/osa:/bin/bash
billy:$1$jwuLPDRU$i6ewGzMTacT3go0gM/Tc11:504:504:~/home/billy:/bin/bash
sammy:$1$dgyjgdssd5gs/jdYJGnc65b/g5.:505:505:~/home/sammy:/bin/bash
tammy:$1$34dfsd23zsjRYJGnc9rPghvb/w.:506:506:~/home/tammy:/bin/bash
ivo:$1$jwuLPDRU$i6cwuzMTacT3go0gM/Tc11:506:506:~/home/ivo:/bin/bash
lenz:$1$FWFvdggJw7bfnc9rP15xgdhd/4.:507:507:~/home/lenz:/bin/bash
lance:$1$dgyjgdssd5gu/jdYJGnc65b/g5.:508:508:~/home/lance:/bin/bash
saul:$1$FWGys23z$oUHswt7bvnc9rP15xvb/w.:509:509:~/home/saul:/bin/bash
tunisia:$1$jwuLPDRU$i6cwGzMTacT3go0gM/Tc11:510:510:~/home/tunisia:/bin/bash
ishmail:$1$jdhfy$d8zMTacT3go0gM/Tc11:511:511:~/home/ishmail:/bin/bash
nevus:$1$FWFvdggJw7bvnc9rP15xgdhd/4.:512:512:~/home/nevus:/bin/bash
kinetic:$1$FWGys23z$oUHJw7bvnc9rP15xvb/w.:513:513:~/home/kinetic:/bin/bash
sasquatch:$1$gJghws23zsjdYJGnc65b/g5.:514:514:~/home/sasquatch:/bin/bash
[root@m10A /root]#

```

Figure 205. *ypcat passwd*

Now to really test the machine, log in to a NIS client using an account that is on the NIS master. When you log in, you should see the following:

```

Connected to netfinity.com.
Escape character is '^]'.
Welcome to SuSE Linux 7.0 (i386) - Kernel 2.2.16 (0).

netfinity login: justin
Password:
Last login: Tue Nov  7 02:36:20 from netfinity.com
Have a lot of fun...
No directory /home/justin!
Logging in with home = "/".
justin@netfinity: / >

```

Figure 206. *No home directory*

Since justin's home directory is on nismaster, we get an error logging in. This can be fixed by creating a home directory for justin on the client if necessary. Another option would be to use NFS in conjunction with NIS to automatically mount justin's home directory.

12.4 Sources of additional information

For further information or troubleshooting, *Managing NFS and NIS* by Hal Stern is good resource. The NIS how-to by Thorsten Kukuk is an excellent place to start. Find it at:

<http://www.metalab.unc.edu/pub/Linux/docs/HOWTO/NIS-HOWTO>.

Chapter 13. LDAP - Lightweight Directory Access Protocol

LDAP has become a buzzword in the IT world. The exciting thing about LDAP and directory services is that they can be used for so many purposes. This chapter will give you a brief explanation of what LDAP is, what it can be used for, basic structures, and simple implementation on the Linux OS. This chapter merely scratches the surface of what is actually possible with LDAP.

13.1 What is LDAP?

LDAP stands for Lightweight Directory Access Protocol. LDAP has become an Internet standard for directory services that run over TCP/IP. LDAP is a client/server protocol for accessing a directory service. Originally designed as a front-end for X.500 databases, LDAP is now commonly used in a stand-alone capacity. IBM, Netscape, Sun, Novell, Microsoft, and many other companies are incorporating LDAP into their directory structures.

13.1.1 Directory Services

A directory service is the collection of software, hardware, processes, policies, and administrative procedures involved in making the information in a directory available to the users of the directory.

A directory is similar to a database. However, directories and databases differ in the number of times they are searched and updated. Directories are tuned for being searched, while relational databases are geared toward maintaining data with a frequent number of updates.

Examples of directories would be the Yellow Pages, a card catalog, or an address book. Information is organized in a defined hierarchy and given attributes.

When we place a directory online, the data becomes dynamic in the sense that it can be easily updated and cross-referenced. Unlike printed material, any updates that occur are instantaneous for all users.

You can apply security to the directory so that only intended users can view, modify, or create data. This security can be based upon groups, individual users, or any other authentication scheme. The data can also be encrypted.

Directory services typically involve data distribution and replication. The advantages of distributing your directory services are performance, availability, and reliability. For a segmented network, distribution of servers containing the directory data improves performance by reducing network

traffic and load on individual servers. By replicating your data on multiple servers you increase availability in case a single server should go down.

13.1.2 X.500

In the mid-1980s, the International Telecommunications Union (ITU, formerly the CCITT) and the International Organization for Standardization (ISO) merged their efforts on directory services standards and created X.500. The X.500 specifications consist of a series of recommendations on the concepts, models, authentication, distribution, attributes, objects, and replication that underlie an X.500 directory service.

Early X.500 implementations used a client access protocol known as DAP. DAP is thick, complicated, and difficult to implement for desktop computers. For all of these reasons other lighter-weight protocols were developed. As predecessors to LDAP, DIXIE and DAS were very successful. Out of this success a group from the Internet Engineering Task Force (IETF) began work on LDAP. The first Request for Comments (RFC 1487) describing LDAP was released in July 1993.

13.2 How can I use LDAP?

LDAP can allow system and network administrators to centrally manage users, groups, devices, and other data. IT decision makers can avoid tying themselves to a single vendor for applications and operating systems. Developers can use LDAP-based standards to ensure cross-platform integration.

Some practical applications of LDAP-based directory services include:

- Corporate address book
- User authentication
- Domain Name System

13.3 LDAP basics

The LDAP information model is based on objects. Objects can be people, printers, servers, or just about anything you can think of. The most basic unit of the LDAP model is the entry. An entry is a collection of information about an object. Each entry belongs to an object class that determines required and optional attributes. Each attribute has a type and one or more values. The type describes the kind of information contained in the attribute and the value contains the actual data.

An LDIF file is the standard way of representing directory data in a textual format. This format can typically be used for importing and exporting directory data. The following is an LDIF file for loading a basic LDAP directory and adding a user that we can use to test our authentication.

```
dn: ou=people, dc=ibm, dc=com
objectclass: top
objectclass: organizationalUnit
description: users who authenticate via the ldap server

dn: uid=ldaptest, ou=People, dc=ibm, dc=com
uid: ldaptest
cn: ldaptest
objectclass: account
objectclass: posixaccount
objectclass: top
objectclass: shadowaccount
userpassword: test123
shadowLastchange:11263
shadowMax: 99999
shadowWarning: 7
loginShell: /bin/bash
uidNumber: 500
gidNumber: 500
homeDirectory: /home/testldap
```

Each LDAP entry must have a distinguished name (DN). The distinguished name is a unique key that refers to that entry specifically.

The first group of entries is to create the root or base entries in the directory. In this case we are using com, ibm, and people (organizational unit).

Within the organizational unit, people we will store the users and their corresponding authentication information. The second group of entries is to add the user ldaptest and the attributes that are necessary for authentication.

13.4 Implementation on Linux

You can download the latest stable version of OpenLDAP from:

```
ftp.openldap.org/pub/openldap
```

However, we will be using the packages that come with the SuSE Linux distribution. Install or verify installation of the following RPMs:

```
openldap
```

```
nss_ldap
pam_ldap
```

13.4.1 Slapd.conf

Now edit the `/etc/openldap/slapd.conf` file. Replace `ibm` and `com` with the name of your organization.

```
# See slapd.conf(5) for details on configuration options.
# This file should NOT be world readable.
#
include /etc/openldap/slapd.at.conf
include /etc/openldap/slapd.oc.conf
schemacheck off
#referral ldap://root.openldap.org/

pidfile /var/run/slapd.pid
argsfile /var/run/slapd.args

#####
# ldbm database definitions
#####

database ldbm
suffix "ou=people, dc=ibm, dc=com"
#suffix "o=My Organization Name, c=US"
rootdn "cn=admin, ou=people, dc=ibm, dc=com"
#rootdn "cn=Manager, o=My Organization Name, c=US"
# cleartext passwords, especially for the rootdn, should
# be avoided. See slapd.conf(5) for details.
rootpw insert_your_password
# database directory
# this directory MUST exist prior to running slapd AND
# should only be accessible by the slapd/tools Mode 700 recommended.
directory /var/lib/ldap
```

Figure 207. `slapd.conf` file

13.4.2 Ldap.conf

Edit the `/etc/openldap/ldap.conf` file. Modify the host and base entries. In this case the LDAP server we will be authenticating against is on the localhost. Replace `padl` with the name of your organization. In our example we used `ibm`. All other entries can be left with the default values.


```

#
# LDAP Defaults
#

# See ldap.conf(5) for details
# This file should be world readable.

BASE ou=people, dc=ibm, dc=com
HOST 127.0.0.1

#HOST ldap.openldap.org ldap-master.openldap.org:666
#PORT 389

#SIZELIMIT 12
#TIMELIMIT 15
#DEREF never

```

Figure 208. *ldap.conf* file

13.4.3 nsswitch.conf

The `nsswitch.conf` file controls which type of name service your host will use to look up various types of entries. With LDAP you can easily set up Domain Name System and host lookup tables. In this example we need to modify the entry for `passwd`. Edit the `/etc/nsswitch.conf` file and add `ldap` to the entry for `passwd` and group lookups:

```

passwd: files, ldap
group: file, ldap

```

13.4.4 /etc/pam.d/login

PAM is a system of libraries that handle the authentication tasks of services on the system. PAM separates the tasks of authentication into four independent management groups: account, authentication, password, and session.

- **account:** provides account verification. Examples: Checking for password expiration or verifying that the user has permission to access the requested service.
- **authentication:** establish the user is who he claims to be, typically by prompting for a password.
- **password:** authentication update mechanism. Example: Prompting the user to enter a new password when the current password has expired.

- **session:** tasks that should be done prior to a service being given and after it is withdrawn. Examples: audit trail maintenance and mounting of the user's home directory.

To enable the PAM modules for accessing the LDAP directory edit the `/etc/pam.d/login` file:

```

#%PAM-1.0
auth    requisite /lib/security/pam_unix.so nullok #set_secrcp
auth    required/lib/security/pam_securetty.so
auth    required /lib/security/pam_nologin.so
auth    sufficient /lib/security/pam_ldap.so use_first_pass
#auth   required /lib/security/pam_homecheck.so
auth    required /lib/security/pam_env.so
auth    required /lib/security/pam_mail.so
account sufficient /lib/security/pam_ldap.so
account required /lib/security/pam_unix.so
password required /lib/security/pam_ldap.so
password required /lib/security/pam_pwcheck.so nullok
password required /lib/security/pam_unix.so nullok use_first_pass use_authtok
session required /lib/security/pam_unix.so none # debug or trace
session required /lib/security/pam_limits.so

```

Figure 209. Login file

For more information on PAM, see the PAM administrator's guide at: <http://www.kernel.org/pub/linux/libs/pam/Linux-PAM-html/pam.html>

13.4.5 Starting OpenLDAP

To start slapd, type `rcldap start`.

With slapd successfully running, we now need to load the initial data.

Create an LDIF file like the one on page 249. Replace `ldaptest` with your user name and `ibm` with your organization name.

Once you have created the `entries.ldif` file, load the LDAP server.

```

/usr/bin/ldapadd -f entries.ldif -D "cn=admin, ou=people, dc=ibm, dc=com"
-w insert_your_password

```

13.4.6 Testing authentication

Now that we've added our test user in the LDAP directory, we need to create the home directory we specified in our LDIF file along with the appropriate ownership.

```
mkdir /home/ldaptest  
chown 500:500 /home/ldaptest
```

Now try to log in as ldaptest with the password test.

If the authentication failed, check the configuration files for typos. It can also be helpful to view the `/var/log/messages` file to see any errors that the LDAP server reports.

Once you have logged in successfully, either remove the test user or at the very least create a properly encrypted password.

13.4.7 Migrating `/etc/passwd`

PADL maintains a group of scripts to allow migration of information to the LDAP server. Using these utilities we can migrate your current user database (`/etc/passwd`) to the LDAP directory. You can find the migration utilities at:

<http://www.padl.com/tools.html>

Download the migration tools, and unzip and untar them:

```
tar -zxvf MigrationTools.tgz
```

In keeping with our example for user authentication, we will look at the script that migrates the entire `/etc/passwd` into an LDIF structure, which can then be easily put into an LDAP directory. Most of the migration scripts are written in perl. Perl must be installed in order to run them.

Change directories to the directory created by the previous command.

```

netfinity: /MigrationTools-27 # ls
.                migrate_all_online.sh
..              migrate_base.pl
CVS VersionInfo.txt migrate_common.ph
Make.rules      migrate_fstab.pl
MigrationTools.spec migrate_group.pl
README          migrate_hosts.pl
migrate_aliases.pl migrate_netgroup.pl
migrate_all_netinfo_offline.sh migrate_netgroup_byhost.pl
migrate_all_netinfo_online.sh migrate_netgroup_byuser.pl
migrate_all_nis_offline.sh migrate_networks.pl
migrate_all_nis_online.sh migrate_passwd.pl
migrate_all_nisplus_offline.sh migrate_protocols.pl
migrate_all_nisplus_online.sh migrate_rpc.pl
migrate_all_offline.sh migrate_services.pl

```

Figure 210. PADL LDAP migration tools

The file `migrate_common.ph` keeps common definitions for all of the migration tools. We need to edit the following lines to be compatible with the directory we've set up:

```

# Default DNS domain
$DEFAULT_MAIL_DOMAIN = "padl.com";
# Default base
$DEFAULT_BASE = "dc=padl,dc=com";

```

PADL is the company who developed the Open Source `nss_ldap` module (see <http://www.padl.com> for more information). Replace `padl` entries with the name of your organization in both the default DNS domain and in the default base.

Stop

The LDIF file created by the migration script will have the encrypted passwords of all the users on the system, including root. Treat this file as you would the `/etc/shadow` file.

Now, as root, run the `migrate_passwd.pl` script against `/etc/passwd`.

```
migrate_passwd.pl /etc/passwd /root/passwd.ldif
```

The final argument is the location and name of the LDIF file we want to create.

Edit the LDIF file and remove any users you do not want added to the LDAP directory.

Now to add the entire contents of the LDIF file, execute the following command:

```
ldapadd -f /root/passwd.ldif -D "cn=admin, ou=people, dc=ibm, dc=com" -w secret
```

Replace the `admin`, `ibm`, and `secret` with the appropriate entries.

After testing, you can remove those users from `/etc/passwd` and `/etc/shadow` and use LDAP to all of your authentication.

In a large environment you can set up your servers to authenticate against a single centralized LDAP directory server or create replicas. Either way, you have just made user administration a lot simpler.

Chapter 14. General performance tools in Linux

Linux offers a great variety of ways to optimize your system for maximum performance. Apart from the general fact that it is always good to have as much RAM and the fastest CPU as possible, there are some additional parameters to tune a Linux system. This section is intended as a collection of useful hints and tools, but without getting into too much detail about them. Please refer to the respective documentation and references. You should also note that using some of these hints may render your system unstable; use them at your own risk and only if you know what you are doing.

14.1 General configuration hints

These are some general tips for tweaking your system to maximize performance.

Recompile your programs and the Linux kernel with all available compiler optimization flags (for example, `-funroll-loops`, `-fomit-frame-pointer`, `-O6`) and all architecture-specific compiler options for your hardware architecture. This may increase the size of binaries or make them unable to run on some processors, but you can gain a lot of speed in comparison with the binaries shipped in the distribution. Alternatively you could use special compilers for your architecture (for example, `pgcc`), which offer even more sophisticated optimization options.

Create swap partitions of equal priority but different hard disk drives to allow load balancing. Please note that they need to be different devices! Using two different partitions on one hard disk will have the reverse effect. Even better, try to avoid swapping at all by adding more memory. A busy server should never need to swap, as this would severely degrade the overall performance.

If you are running a heavily loaded server with a lot of parallel processes, you might run into the Linux kernel's limit of running processes (512 by default). This maximum number of tasks is configurable in the kernel sources, so you have to recompile the kernel after changing this value. This value is defined in the file `/usr/src/linux/include/linux/tasks.h`:

```
#define NR_TASKS      512
```

You can increase this value up to 4090 processes, if necessary.

Linux offers a file system mount option that is called `noatime`. The `atime` is a timestamp of the last access time (reading and writing) for a certain file. This option can be added to the mount options in the `/etc/fstab` file. When a file

system is mounted with this option, read accesses to files will no longer result in an update of the inode access time information. This information is usually not very interesting on a file or Web server, so the lack of updates to this field is not relevant. The performance advantage of the `noatime` flag is that it suppresses write operations to the filesystem for files that are simply being read. Since these write accesses add additional overhead, this can result in measurable performance gains. Instead of specifying this as a mount option that would apply to the whole file system, you can use the command `chattr` to set this flag on single files or directories. For example:

```
chattr -R +A /var/spool/news
```

This command would set the `noatime` flag recursively on all files below the news spool directory (a very common practice on busy news servers). See the manual page `chattr(1)` for more information.

You can use the `hdparm` tool to tune some hard disk drive parameters. Unfortunately most of them only work on IDE systems (which should be avoided in server systems, anyway), but the option `-a` works for SCSI, too. The manual page describes it as follows: “This option is used to get/set the sector count for file system read-ahead. This is used to improve performance in sequential reads of large files, by prefetching additional blocks in anticipation of them being needed by the running task. The default setting is 8 sectors (4 KB). This value seems good for most purposes, but in a system where most file accesses are random seeks, a smaller setting might provide better performance. Also, many drives have a separate built-in read-ahead function, which alleviates the need for a file system read-ahead in many situations.” For example, to set the sector count read-ahead of your first SCSI disk to 4 sectors (2 KB), you would use the following command:

```
hdparm -a 4 /dev/sda
```

See the `hdparm` manual page for a complete list of available options.

14.1.1 Powertweak

`Powertweak`, by Dave Jones of SuSE Labs, is a general-purpose application to allow you to “tweak” different system parameters.

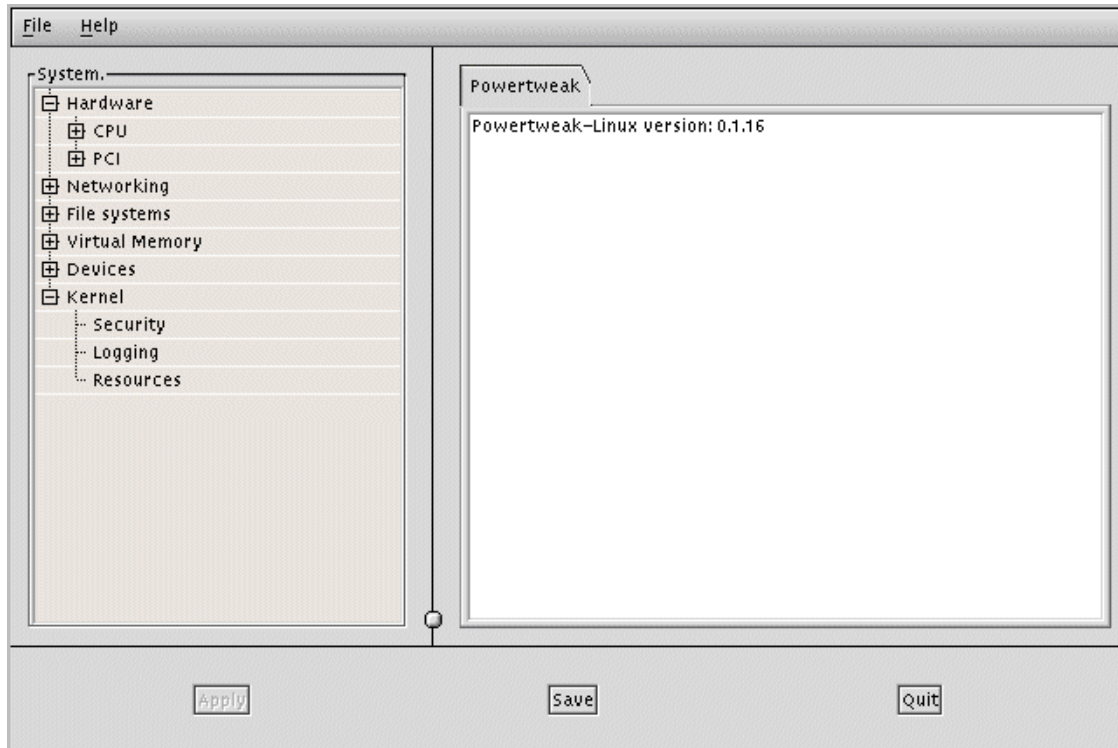


Figure 211. The main Powertweak window

As you can see in Figure 211, Powertweak allows tweaking of many system parameters. Some tabs will only display information about a device and will not allow you to tweak it. This is either because the setting for the device cannot be manipulated, or support has not been integrated with Powertweak yet.




Figure 212. Powertweak: CPU information

The information about a device or a system resource is usually very detailed. This is ideal for troubleshooting the system, or reporting device IDs to vendors and so on.

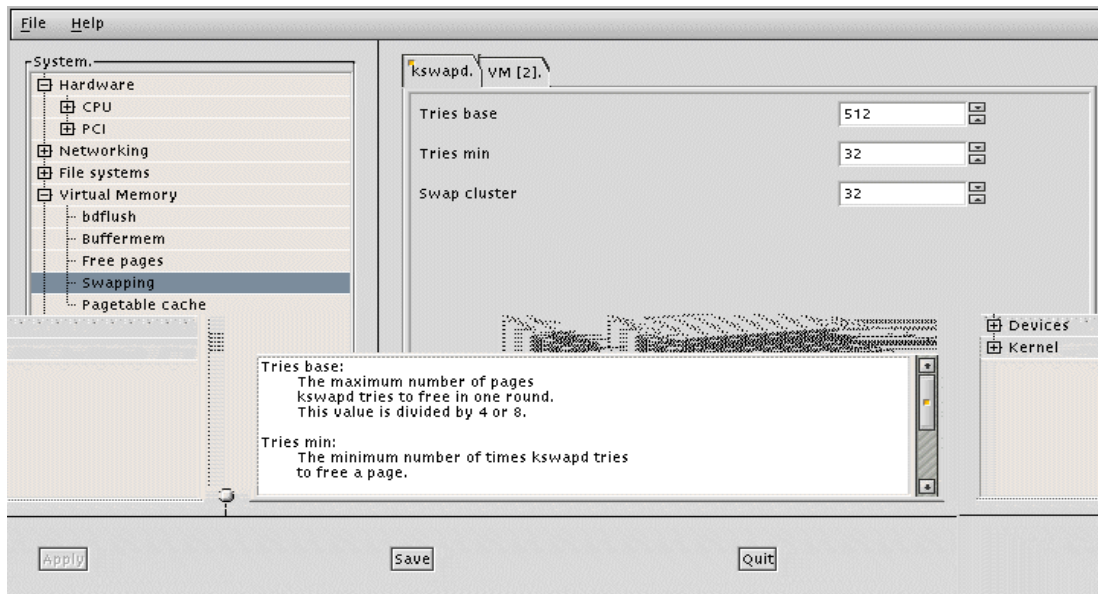


Figure 213. Powertweak swapping parameters

Details of what the settings do to is given for every dynamic resource that PowerTweak can manipulate. As with all tweaking it is a matter of trial and error. As such, do not try to manipulate the system in this way on a production/running system. It is advisable to test it out on a system of the same specification and then use the optimal settings as a basis for the production system.

14.1.2 Services

You should disable all unused services and daemons, especially network-related services. This has several advantages: fewer open services need fewer system resources (file descriptors, memory) and the system is less vulnerable to external attacks against known security holes. A good starting point is the `/etc/inetd.conf` file. Comment out all services you do not need, or disable `inetd` completely. This can also be done using `Yast2`.

The Linux `/proc` filesystem offers a lot of entry points for run-time optimization without recompiling the kernel. This directory does not physically exist on your hard drive; it is mapped as a virtual directory. Most of the files contained there are readable and contain various system information. Other files can be edited with a regular text editor to set a certain kernel parameter. See `/usr/src/linux/Documentation/sysctl/README` in the Linux kernel sources for

a detailed description of the tunable parameters (including file system, virtual memory, etc.).

There are some special TCP options that can be disabled in a local network with high signal quality and bandwidth, since they are mostly intended for loss connections (see `/usr/src/linux/net/TUNABLE` in the Linux kernel sources for a detailed list):

To disable TCP timestamps, enter:

```
echo 0 > /proc/sys/net/ipv4/tcp_timestamps
```

To disable window scaling, enter:

```
echo 0 > /proc/sys/net/ipv4/tcp_window_scaling
```

To disable selective acknowledgments, enter:

```
echo 0 > /proc/sys/net/ipv4/tcp_sack
```

To tune the default and maximum window size (only if you know what you are doing), enter:

```
/proc/sys/net/core/rmem_default - default receive window  
/proc/sys/net/core/rmem_max     - maximum receive window  
/proc/sys/net/core/wmem_default - default send window  
/proc/sys/net/core/wmem_max     - maximum send window
```

The following Web sites offer a lot of additional helpful hints about tuning and performance issues on Linux:

```
http://tune.linux.com  
http://www.tunelinux.com  
http://linuxperf.nl.linux.org/
```

14.1.3 Kernel recompilation

Recompiling the kernel to include only the drivers and features needed by a machine can help to decrease the amount of memory used in the system.

To be able to recompile the kernel, you will need the C development system installed on the system. The development environment can be found in the `d (devel)` package section. You will also need the SuSE kernel source and headers, these can be found as the packages, `lx_suse` and `linclude`.

Once you have installed the development system, and the kernel source code, you have a choice of how to set up the kernel for compilation:

- **make config** - A series of questions are asked about settings and drivers etc you wish to compile into the kernel (console based).
- **make menuconfig** - A menu-based tool to set up the kernel compilation. Based on ncurses (console based).
- **make xconfig** - an X application to configure the kernel compilation. You will need to be running X-Windows to use this version.

change to the `/usr/src/linux-2.2.16.SuSE` directory and type:

```
make config_type
```

where `config_type` is one of the above settings.

Once you have set up the compilation environment in the same directory, type the following, pressing Enter after each:

```
make clean
make dep
make bzImage
```

The resulting kernel image will be stored as `/usr/src/linux/2.2.16.SuSE/arch/i386/boot/bzImage`.

If you compiled any features as modules, in the `/usr/src/linux-2.2.16.SuSE` directory you will need to type:

```
make modules
make modules_install
```

You will need to copy the `bzImage` file to the `/boot` directory and use YaST to install the image as a new kernel. See section 3.6.2, “Kernel and boot configuration” in the SuSE manual.

Here are a few guidelines to follow when selecting the drivers and features to be used in a kernel:

- Drivers that are needed constantly by the server should be compiled directly into the kernel.
- Drivers that are needed by the system, but will not be in constant use, should be compiled as modules. This could be the case for an IDE CD-ROM drive.
- You are given the opportunity to set default values for a number of system resources, including timeouts. Set these to realistic values that can reduce times for device/resource access. Make sure you understand what you are changing with these values. You can usually look in the source code of the relevant driver to view comments about certain settings.

- Do not enable resources in the kernel that are not essential to the system. This includes framebuffer and sound support. These are nice things to have, but are not essential to the system that will be used as a server.
- If a new driver for a resource or device becomes available, use it in the kernel. You are usually given instructions on how to integrate these drivers into your system.
- With regards to the above comment, it is essential you upgrade to the correct ServeRAID driver when you upgrade the firmware, not only for speed, but for the integrity of your system.

14.2 System monitoring and performance test tools

This section introduces a small collection of useful tools, among the many available, to monitor your Linux system or to gather system information.

To get an overview about all running processes and the system load, run the command `top` in a terminal session.

```

11:53am up 3:57, 1 user, load average: 0.00, 0.00, 0.00
34 processes: 33 sleeping, 1 running, 0 zombie, 0 stopped
CPU states: 0.0% user, 1.6% system, 0.0% nice, 98.4% idle
Mem: 62968K av, 59196K used, 3772K free, 17408K shrd, 15164K buff
Swap: 125996K av, 0K used, 125996K free 33768K cached

```

PID	USER	PRI	NI	SIZE	RSS	SHARE	STAT	LIB	%CPU	%MEM	TIME	COMMAND
515	root	20	0	792	792	628	R	0	1.6	1.2	0:01	top
1	root	0	0	196	196	168	S	0	0.0	0.3	0:04	init
2	root	0	0	0	0	0	SW	0	0.0	0.0	0:00	kflushd
3	root	0	0	0	0	0	SW	0	0.0	0.0	0:00	kupdate
4	root	0	0	0	0	0	SW	0	0.0	0.0	0:00	kpiod
5	root	0	0	0	0	0	SW	0	0.0	0.0	0:00	kswapd
6	root	0	0	0	0	0	SW	0	0.0	0.0	0:00	md_thread
76	root	0	0	648	648	536	S	0	0.0	1.0	0:00	syslogd
79	root	0	0	816	816	392	S	0	0.0	1.2	0:00	klogd
116	at	0	0	552	552	456	S	0	0.0	0.8	0:00	atd
121	root	0	0	452	452	376	S	0	0.0	0.7	0:00	gpm
132	root	0	0	1592	1592	1488	S	0	0.0	2.5	0:00	httpd
135	root	0	0	624	624	528	S	0	0.0	0.9	0:00	lpd
137	wwwrun	0	0	1592	1592	1500	S	0	0.0	2.5	0:00	httpd
138	wwwrun	0	0	1592	1592	1500	S	0	0.0	2.5	0:00	httpd
139	wwwrun	0	0	1592	1592	1500	S	0	0.0	2.5	0:00	httpd
140	wwwrun	0	0	1592	1592	1500	S	0	0.0	2.5	0:00	httpd
141	wwwrun	0	0	1592	1592	1500	S	0	0.0	2.5	0:00	httpd

Figure 214. Example output of `top`

`Top` updates the process list at regular intervals. Press “?” to get an online help screen about the available parameters. To change the refresh interval, press “s” and enter the desired number of seconds between each update. If

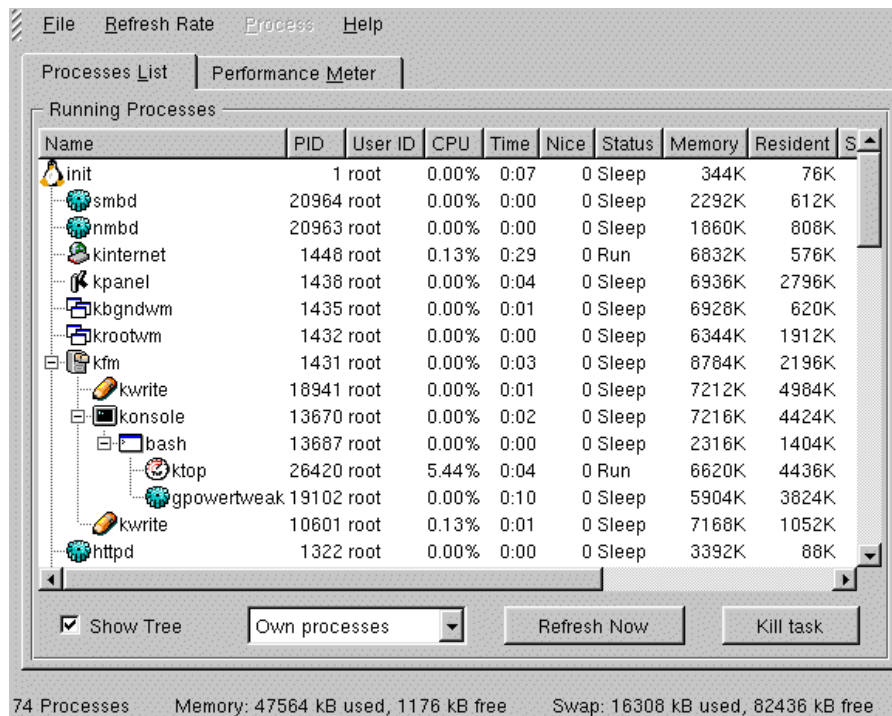


Figure 216. KDE Task Manager: Processes List window

KTop offers two different views. It can either display a list of processes (similar to `top` and `pstree`), or you can switch to the performance meter, which displays the system load and memory usage over a longer time period.




Figure 217. KDE Task Manager: Performance Meter

The KDE control center also gives you a lot of information about your system by reading a number of informative files in the `/proc` filesystem. They can also be displayed in a regular text viewer (for example `more`, `less` or `cat`).

The `/proc/cpuinfo` file contains information about your CPU (that is, vendor, MHz, and flags). For example:

```

SuSE:~ # cat /proc/cpuinfo
processor       : 0
vendor_id     : GenuineIntel
cpu family    : 6
model         : 5
model name    : Pentium II (Deschutes)
stepping      : 2
cpu MHz       : 513.953346
cache size   : 512 KB
fdiv_bug     : no
hlt_bug      : no
sep_bug      : no
f00f_bug     : no
coma_bug     : no
fpu          : yes
fpu_exception : yes
cpuid level   : 2
wp           : yes
flags        : fpu vme de pse tsc msr pae mce cx8 sep mtrr pge mca cmov pat p
se36 mmx osfxsr
bogomips     : 313.75

```

Figure 218. Cpuinfo file

The `/proc/interrupts` file lists all interrupts used by Linux. Note that this shows interrupts only from devices that have been detected by the kernel! If a device will not be detected because of a resource conflict, you have to resolve this conflict manually (for example, by changing the BIOS setup). For example:

```

SuSE:~ # cat /proc/interrupts
CPU0
0: 548029 XT-PIC timer
1: 557 XT-PIC keyboard
2: 0 XT-PIC cascade
8: 2 XT-PIC rtc
9: 371 XT-PIC PCnet/PCI II 79C970A
12: 68 XT-PIC PS/2 Mouse
13: 0 XT-PIC fpu
14: 198235 XT-PIC ide0
15: 3 XT-PIC ide1
NMI: 0

```

Figure 219. Interrupts file

The `/proc/ioports` file contains all allocated device I/O ports. The same note as for interrupts applies here. Only devices that are actually detected by the kernel are listed here. For example:

```

SuSE:~ # cat /proc/ioproports
0000-001f : dma1
0020-003f : pic1
0040-005f : timer
0060-006f : keyboard
0070-007f : rtc
0080-008f : dma page reg
00a0-00bf : pic2
00c0-00df : dma2
00f0-00ff : fpu
0170-0177 : ide1
01f0-01f7 : ide0
02e8-02ef : serial(auto)
02f8-02ff : serial(auto)
0376-0376 : ide1
03c0-03df : vga+
03e8-03ef : serial(auto)
03f6-03f6 : ide0
03f8-03ff : serial(auto)
1000-101f : PCnet/PCI II 79C970A
1020-1027 : ide0
1028-102f : ide1

```

Figure 220. *ioports* file

The `/proc/meminfo` file displays information about memory (for example, memory used, free, swap size). You can also use the `free` command to display this information. For example:

```

SuSE:~ # cat /proc/meminfo
      total:      used:      free:  shared: buffers:  cached:
Mem:  64569344  62578688  1990656  54308864  18792448  27807744
Swap: 129019904   102400 128917504
MemTotal:        63056 kB
MemFree:         1944 kB
MemShared:       53036 kB
Buffers:         18352 kB
Cached:          27156 kB
SwapTotal:      125996 kB
SwapFree:       125896 kB
SuSE:~ # free
      total      used      free      shared    buffers    cached
Mem:        63056    61124    1932     53068     18352     27164
-/+ buffers/cache:    15608    47448
Swap:      125996      100    125896

```

Figure 221. *Meminfo* file

The `/proc/mounts` file shows all currently mounted partitions. The `mount` command without parameters will display similar information. For example:

```

SuSE:~ # cat /proc/mounts
/dev/root / ext2 rw 0 0
proc /proc proc rw 0 0
/dev/hda1 /boot ext2 rw 0 0
devpts /dev/pts devpts rw 0 0
SuSE:~ # mount
/dev/hda3 on / type ext2 (rw)
proc on /proc type proc (rw)
/dev/hda1 on /boot type ext2 (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=0620)

```

Figure 222. Mounts file

The `/proc/partitions` file displays all existing partitions on all devices. You can also use `fdisk -l` to display this information. For example:

```

SuSE:~ # cat /proc/partitions
major minor #blocks name

 3      0  1023907 hda
 3      1     6016 hda1
 3      2  126000 hda2
 3      3  891072 hda3
 3     64  1023907 hdb
 3     65  1023088 hdb1
22      0 1073741823 hdc
SuSE:~ # fdisk -l

Disk /dev/hda: 32 heads, 63 sectors, 1015 cylinders
Units = cylinders of 2016 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/hda1 *            1           6        6016+    83  Linux
/dev/hda2              7          131       126000    82  Linux swap
/dev/hda3             132         1015       891072    83  Linux

Disk /dev/hdb: 32 heads, 63 sectors, 1015 cylinders
Units = cylinders of 2016 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/hdb1              1          1015     1023088+    83  Linux

```

Figure 223. Partition file

The `/proc/pci` file gives information about all your PCI devices. You can also use the `lspci` command to provide output that is easier to read. Please note that `/proc/pci` is obsolete and will be replaced by `/proc/bus/pci/*` in the future. For example:

```

bash-2.04# cat /proc/pci
PCI devices found:
  Bus 0, device 0, function 0:
    Host bridge: Intel 82439TX (rev 1).
    Medium devsel. Master Capable. No bursts.
  Bus 0, device 7, function 0:
    ISA bridge: Intel 82371AB PIIX4 ISA (rev 8).
    Medium devsel. Master Capable. No bursts.
  Bus 0, device 7, function 1:
    IDE interface: Intel 82371AB PIIX4 IDE (rev 1).
    Medium devsel. Fast back-to-back capable. Master Capable. Latency=64.
    I/O at 0x1020 [0x1021].
  Bus 0, device 15, function 0:
    Display controller: Unknown vendor Unknown device (rev 0).
    Vendor id=15ad. Device id=710.
    Medium devsel. Fast back-to-back capable. Master Capable. Latency=64.
    I/O at 0x1030 [0x1031].
    Non-prefetchable 32 bit memory at 0xfc000000 [0xfc000000].
    Non-prefetchable 32 bit memory at 0xfb000000 [0xfb000000].
  Bus 0, device 16, function 0:
    Ethernet controller: AMD 79C970 (rev 16).
    Medium devsel. Fast back-to-back capable. IRQ 9. Master Capable. Latency=64.
    Min Gnt=6.Max Lat=255.
    I/O at 0x1000 [0x1001].
    Non-prefetchable 32 bit memory at 0xfd000000 [0xfd000000].
bash-2.04# lspci
00:00.0 Host bridge: Intel Corporation 430TX - 82439TX MIMC (rev 01)
00:07.0 ISA bridge: Intel Corporation 82371AB PIIX4 ISA (rev 08)
00:07.1 IDE interface: Intel Corporation 82371AB PIIX4 IDE (rev 01)
00:0f.0 Display controller: VMWare Inc Virtual SVGA
00:10.0 Ethernet controller: Advanced Micro Devices [AMD] 79c970 [PCnet LANCE] (rev 10)

```

Figure 224. PCI file

The `/proc/swaps` file displays information about all active swap partitions. For example:

```

bash-2.04# cat /proc/swaps
Filename                                Type              Size    Used    Priority
/dev/hda2                               partition         125996  56      -1

```

Figure 225. Swaps file

The `/proc/version` file displays some version information about the Linux kernel. The command `uname -a` will display similar information. For example:

```
bash-2.04# cat /proc/version
Linux version 2.2.16 (root@Pentium.suse.de) (gcc version 2.95.2 19991024 (release)) #1
Wed Aug 2 20:22:26 GMT 2000
bash-2.04# uname -a
Linux netfinity 2.2.16 #1 Wed Aug 2 20:22:26 GMT 2000 i686 unknown
```

Figure 226. Version file

If you want to obtain some more information about your SCSI devices, have a look at the files below `/proc/scsi`.

A tool that is also gathering system information from the `/proc` file system is `vmstat`. It reports information about processes, memory, paging, block IO, traps, and CPU activity. The first report produced gives averages since the last reboot. Additional reports give information on a sampling period of length delay. The process and memory reports are instantaneous in either case. `vmstat` is very helpful for logging CPU and memory usage over a longer period of time.

Apart from configuring numerous parameters of your hard drive, the command `hdparm` can also be used to perform hard disk performance tests with the command `hdparm -tT <device>`. For example:

```
SuSE:~ # hdparm -tT /dev/hda

/dev/hda:
Timing buffer-cache reads: 64 MB in 0.68 seconds =94.12 MB/sec
Timing buffered disk reads: 32 MB in 29.51 seconds = 1.08 MB/
SuSE:~ # hdparm -c1 /dev/hda

/dev/hda:
setting 32-bit I/O support flag to 1
I/O support = 1 (32-bit)
SuSE:~ # hdparm -tT /dev/hda

/dev/hda:
Timing buffer-cache reads: 64 MB in 0.67 seconds =95.52 MB/sec
Timing buffered disk reads: 32 MB in 12.92 seconds = 2.48 MB/sec
```

Figure 227. Hdparm performance test

Another popular hard disk performance test is `bonnie`, found at <http://www.textuality.com/bonnie/> (an RPM package for SuSE Linux is included in the distribution). Note, however, that these tests are mostly useful for testing different parameter settings on one machine as a relative measure, not as a comparison between different systems.

To test the throughput of your network, you can either use `netperf`, found at <http://www.netperf.org/netperf/NetperfPage.html> or `bing` (included in SuSE Linux).

Chapter 15. Setting up a Beowulf cluster

For a long time, parallel computing has been a domain of commercial vendors. By using Linux, it is now possible to create a powerful supercomputer using regular PCs with off-the-shelf components that are networked together with fast Ethernet cards or by using special high-speed interconnections such as SCI or Myrinet. Beowulf clusters offer high performance computing at a fraction of the cost of a regular parallel computer (the price/performance ratio is usually between three and ten times better than for a “regular” supercomputer).

Beowulf was “the son of Scyld in the Scandian lands”, a character from one of the oldest English epic poems. The legend tells that he defeated a monster called “Grendel” (see <http://legends.dm.net/beowulf/index.html> and <http://www.lnstar.com/literature/beowulf/beowulf.html> for historical background).

The first Beowulf cluster was set up by Donald Becker and Tom Sterling at the NASA Goddard Space Flight Center in 1994. Don Becker is also well known in the Linux community for his work on network drivers in the Linux kernel.

To make use of the parallelism, your software needs to be distributable between the nodes of a cluster. One way is to use libraries such as PVM (Parallel Virtual Machine) or LAM/MPI (Local Area Metacomputer/Message Passing Interface). Regular programs are not suitable for distributed computing. There is now a special load-sharing software called MOSIX, which allows transparent process migration in a cluster. MOSIX can be used by any software that spawns multiple processes or threads. MOSIX requires a special patched Linux kernel and will not be covered here.

This chapter focus as on how to set up PVM and how to demonstrate the parallel computing power using a special version of the famous raytracing software POVray, called PVMPOVray. Raytracing is a method to create realistic images of a scene that is only described by coordinates, light sources, textures, and surface properties such as reflectivity or opacity. The raytracer now computes the reflections, shadows, and refractions of all light rays in the picture and generates the respective image of this scene. Further information about POVray can be found on the POVray Web site at <http://www.povray.org>.

To set up a simple Beowulf cluster, you need at least two PCs running Linux and a functional TCP/IP network connection between them. Regular Ethernet is fine for starters; however, it does not offer the best performance, since it

has a rather high latency, which is crucial if you run applications that need to communicate a lot between nodes. To enable the communication between the nodes, PVM needs to be installed on all these machines as well. XPVM is a useful tool to monitor the communication and setup of the virtual machine, if the number of nodes is not too high (approximately 20-30). XPVM only needs to be installed on the master server.

When using SuSE Linux, make sure that the following packages are installed on all machines in the cluster. These packages can be found in the “beo” package series:

- pvm
- povray
- pvmpov

One machine acts as the master node that distributes jobs to the “slave” nodes. They should share a common work directory (NFS) and it should be possible to run a remote shell `rsh` from each node to another without being prompted for a password (edit the `/etc/hosts.equiv` file on each machine or create a `~/.rhosts` file in the home directory of the user who wants to spawn jobs on remote machines). Start the PVM console by typing `pvm` on the command line. At the PVM command prompt `pvm>`, use the command `add <Hostname>` to add nodes to your virtual machine. PVM now attempts to start the PVM daemon process on the remote machine using `rsh`. If this fails, have a look at the log files on the remote machine. The command `conf` gives you a list of all nodes in your cluster that have successfully been added to PVM. Use `quit` to return to the shell. Alternatively, you can create a file that contains the names of all hosts that you want to use for your cluster (one on each line) and run `pvm <hostfile>`. This will automatically add all these hosts to the virtual machine. This is basically all you need to set up a basic Beowulf cluster. To make use of the parallel computing power, you now need to have a program that has been written using the PVM library. One example here is PVMPOVray.

To run XPVM, you first have to set the following environment variable:

```
export XPVM_ROOT=/usr/X11R6/lib/xpvm/
```

Now you can start `xpvm` by typing `xpvm` in a terminal window. Add the other nodes by clicking **Hosts...** -> **Other Hosts...** An icon should appear for each host that has been successfully added to the virtual machine. Click **Tasks...** -> **SPAWN** to start the distribution of a job. To give a demonstration, spawn the following command:

```
/usr/X11R6/bin/x-pvmpov +I/usr/lib/povray3/include  
+I/usr/lib/povray3/povscn/level2/skyvase.pov +O skyvase.tga +D +W640  
+H400 +N
```

Set NTasks to the number of hosts involved.

A window should now pop up, and the picture will be created tile by tile. The finished image can be found as “skyvase.tga” in your home directory.

Chapter 16. Backup and recovery

It may seem obvious that backing up and restoring data quickly is critical, but many administrators leave this task at the end of the “to do” list until it is too late. With the ease of use of the commercially available packages BRU (Enhanced Software Technologies), BackupEDGE/RecoverEDGE (MicroLite) or Arkeia (Knox Software), there is no need to wait.

Note

We recommend that you do not connect tape devices to the IBM ServeRAID adapter. Use a separate SCSI controller for the tape devices.

16.1 BRU

BRU is a backup and restore utility with significant enhancements over other common utilities such as tar, cpio, volcopy and dump. BRU is designed to work with most backup devices, including cartridge, 4mm DAT, 8mm (Exabyte) and 9-track tape drives.

BRU includes incremental backups, full backups, multivolume archives, distribution and updates, error detection and recovery, random access capabilities, file comparisons, file overwrite protection, and increased speed over previous versions.

16.1.1 Installing BRU

Before you begin, you need to know the following:

1. The device name of your tape drive. Typically under SuSE Linux this will be `/dev/st0` for the rewinding and `/dev/nst0` for the non-rewinding drive.
2. The size of your backup medium in megabytes.

To install BRU from the floppy drive with the `tar` command, type:

```
cd /tmp
tar xvf /dev/fd0
./install
```

Follow the prompts regarding readme files and licenses, enter your *license data* and your *BRU serial number* when asked to do so until you come to the following window:

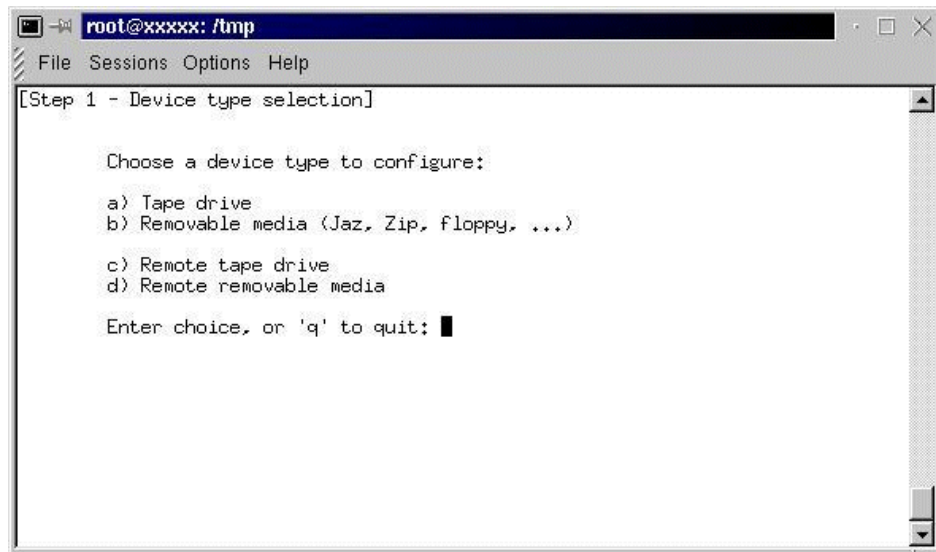


Figure 228. Selecting your backup devices

Enter the letter for your backup device and answer the following questions appropriate for your device.

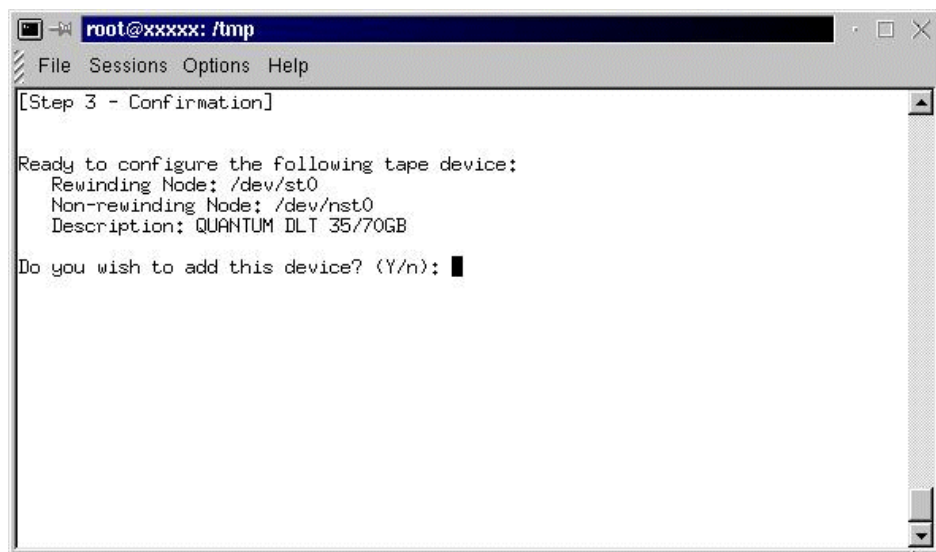


Figure 229. You have entered your backup devices

If you have entered the information for all your backup devices, you will be asked if you would like to install the X11 interface. Select **Y**.

The installation program needs to create an xbru directory. You can select a path or accept the default `/usr/local/`.

The installation program will install executables in a user-specified directory. The default is `/usr/local/bin`.

Note

The key configuration file is `/etc/brutab`. Consult the *BRU User's Guide* for advanced information. Do not edit unless you know what you are doing.

BRU is now installed.

16.1.2 Basic commands

The basic command structure for BRU is:

```
# bru modes [control options] [selection options] [files]
```

Where `bru` is the command or program followed by the mode specifying backup, restore, or various queries. `Control options` specify devices and buffer size. `Selection options` control which files or directories to work with. `Files` is the specified target of the `bru` command.

16.1.3 Basic backup

To back up a single file `/home/ayne/.profile`:

```
# bru -c -vvvv -G /home/ayne/.profile
```

To back up the complete directory `/home/ayne`:

```
# bru -c -vvvv -G /home/ayne
```

To back up the entire system:

```
# bru -c -vvvv -G /
```

16.1.4 Basic restore

To restore a single file `/home/ayne/.profile`:

```
# bru -x -vvvv -ua -w /home/ayne/.profile
```

To restore the complete directory `/home/ayne`:

```
# bru -x -vvvv -ua -w /home/ayne
```

To restore the entire system:

```
# bru -x -vvvv -ua -w /
```

16.1.5 Basic verification and listing commands

The `-i` mode can be used in conjunction with a backup command or by itself. The `-i` mode reads each block of data and verifies the checksum of the block. If used with the verbosity options (`-vvvv`), BRU will give a complete listing of the contents of an archive.

The `-G` mode displays the archive header block, which contains detailed information on the archive including the command used to create the archive. See the *BRU User's Guide* for more information.

The `-gg` mode displays the contents of the on-tape directory. This mode can only be used if the archive was created with the `-G` option.

16.1.6 X Interface

To use BRU's X interface, you will need to be in an X-Windows environment.
Type:

```
xbrn
```

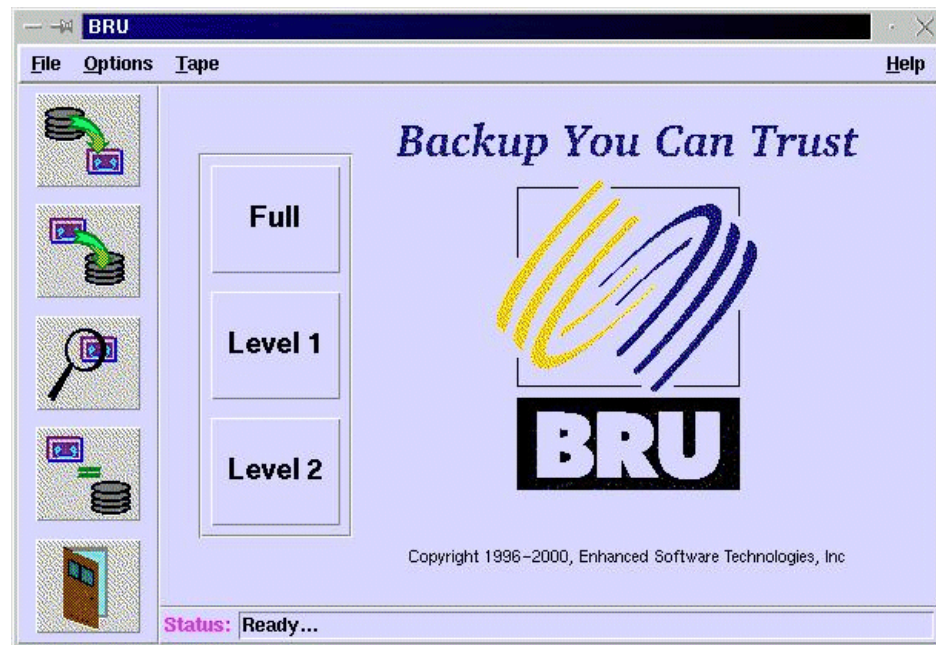


Figure 230. XBRU window

You will see a window similar to Figure 230.

From this interface you can:

- Create and restore backups.
- Create save, and load backup definitions.
- Schedule backups.
- List and verify the contents of archives.
- View the BRU log.

16.1.7 The big buttons in BRU

The three main buttons (Full, Level 1, and Level 2) are shortcuts to various levels of backing up your system, directories, or individual files.

- Select **Full** to back up all the files in the user's home directory, or, if the user is root, the entire system.

- Select **Level 1** to execute a backup for the same files as listed above, on the condition that files have been modified since the previous full backup. If no previous full backup has been done, this will be considered a full backup.
- Select **Level 2** to execute a backup for the same files as listed above, on the condition that files have been modified since the previous level 1 backup. If no previous level 1 backup has been done, this will be considered a level 1 backup.

16.1.8 Creating archives

Creating archives with BRU's X interface is simple. Click the **Backup** button to bring up the Backup File Selection interface (Figure 231).




Figure 231. Creating an archive

The box on the left displays the contents of the current directory (CD:). You can change the current directory by editing the CD entry. Then press Enter.

You can add or remove files and directories from the backup list by selecting them and clicking the appropriate button.

BRU also provides a search function. Click the **Search** button to bring up a dialog box prompting you for a search string. This string can contain typical wildcards.

Backup Definitions are a way to define a set of commonly used backup options or preferences for use at a future time. You can create definitions for use with the backup scheduler or simply use the default selections.

After you have selected the files and directories that you wish to back up, you can click the **Options** button. In this dialog (Figure 232) you can set your preferences regarding different options. After you have made your decisions, click the **Close** button to return to the previous dialog. To start the backup click the **Start Backup** button.



Figure 232. Dialog for backup options

Enter in the next dialog, click **Enter Archive Label** and enter text to identify your new archive. Click **Create Backup** to proceed.

The backup will inform you of how many directories/files and which amount of data will be backed up. During backup, you see a window, informing you about the progress and the actual action. When the backup process has finished, click **Done** to return to XBRU's main dialog.

16.1.9 Scheduling

To access the scheduling feature, go to **File>Scheduler** on the menu.

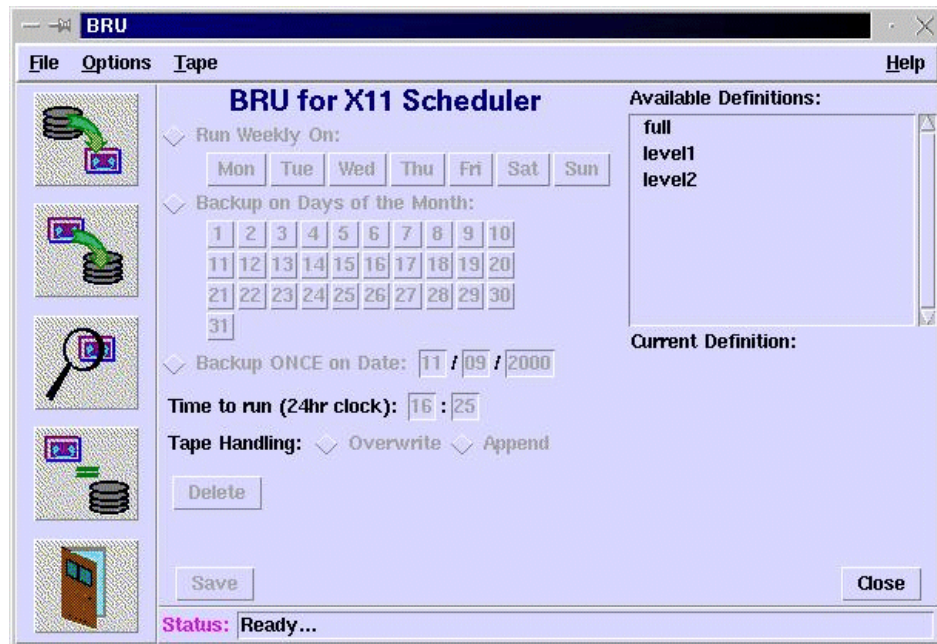


Figure 233. Scheduler

BRU provides a scheduling utility to automate the backup process for the busy administrator. There are three predefined definitions: Full, Level 1, and Level 2. These are the same definitions used in 16.1.7, “The big buttons in BRU” on page 283. You can create your own definitions in the Creating Archives interface.

From the BRU for X11 Scheduler interface, you can set scheduled backups based on weekly, monthly, or single dates. The scheduler is very flexible. In order to take advantage of the scheduling options, you must save your desired schedule configuration and verify that the scheduler is being run from cron. To verify or add the cron entry, log in as root and type:

```
crontab -e
```

Insert the following line:

```
0/5 * * * * /usr/local/bin/bruschedule
```

If you chose a different path for the binaries during installation, change the entry accordingly.

Save the crontab entry. You can now schedule backups.

16.1.10 Restoring files

Restoring files with BRU's X interface is simple. BRU will retrieve the contents of the archive when you click the **Restore** button. After scanning the archive, the Restore File Selection interface (similar to Figure 231) will appear.

Note

If the on-tape directory is not in the archive, then BRU must scan the entire archive to get a listing. This can be very time consuming. When creating an archive, use the -G option to create the on-tape directory or chose **Create On-Tape Directory** in XBRU's **Options** dialog from the backup dialog.

The box on the left displays the contents of the current directory that is stored on the tape. You can change the current directory by editing the **CD:** entry and pressing Enter.

You can add or remove files and directories from the backup list by selecting them and selecting the appropriate button.

When you have selected all of the files and directories that you wish to restore, click the **Restore** button. A progress window will show each file as it is restored.

16.1.11 Listing and verifying archives

For listing the contents of an archive, BRU gives you three options:

1. Header - This option shows the archive header record, which lists the label, creation date, version, and serial number. For more information on the header, consult the *BRU User's Guide*.
2. Filenames only - This option displays the on-tape directory. If the archive was created without using the -G option, BRU will scan the entire archive to create a list of files. You will be prompted before this occurs, as this can be a lengthy process.
3. Full details - This option scans the entire archive for details such as file names, permissions, owners, size, modification times, etc. This process can be time consuming.

For verifying archives, BRU gives you two options:

1. Checksum Verification - When archives are written, a checksum is calculated for each block of data. The checksum is stored in the header of each block. Checksum verification will read each block, recalculate the checksum, and compare the checksum to the value in the header. Each file will be listed as it is verified, along with any errors found. If no errors are found, you know you have an accurate backup.
2. Compare Verification - BRU compares the files in the archive to the files on the hard drive. Any differences, such as modification times, size, or files in the archive that are nonexistent on the hard drive are noted. An *end of differences* notice will be posted when the verification is complete.

16.1.12 Summary

For information on advanced features consult your *BRU User's Guide* or the BRU Web site at:

<http://www.estinc.com/>

16.2 Microlite BackupEDGE

BackupEDGE is a complete backup solution for the Linux platform. It is easy to use and still very robust. With BackupEDGE you can safely archive every file, directory, device node and special file on your file systems. Unlike the standard UNIX tar command, which ignores many important files, BackupEDGE also verifies every byte of data written to the tape to ensure the tape is an accurate reflection of your data. Below are the features provided by BackupEDGE backup software:

- Data Compression - automatic data compression is supported.
- Menu Interface - almost all functions can be accessed through an intuitive menu system.
- Remote Tape Drive Support - you can back up computers across the network.
- High Performance - advanced double buffering and variable block factors.
- Virtual File Support - you can back up virtual (sparse) files.
- Multi-Volume / Multi-Device Archives - automatic spanning across multiple volumes or devices.
- Wildcard Support - when selecting files you can use a wildcard.
- Raw Device Backups - you can archive an entire raw device/partition to tape.

- Master / Incremental Backups
- Unattended Operation - you can perform a master backup or back up only the changed files.

BackupEDGE is designed to operate on Linux kernels 2.x and there are available versions for several types of libraries.

In the following sections we describe how to install, configure and use the Microlite BackupEDGE backup software.

Note

We recommend that you do not connect tape devices to the IBM ServeRAID adapter. Use a separate SCSI controller for the tape devices.

16.2.1 Installing Microlite BackupEDGE

Before you install BackupEDGE you must identify the device entry for your backup device. Usually tape devices under Linux are assigned in device nodes `/dev/st0`, `/dev/st1`... A no-rewind device is created for each tape device, which is `/dev/nst0`, `/dev/nst1`... In our example, we used `/dev/st0` as tape device and `/dev/nst1` as the no-rewind device.

In our example, we used diskette as the installation medium. To install the product, follow these steps:

1. Log in as root.
2. Change the directory to root `/`.
3. Insert the diskette with the product in the floppy drive and execute the command:

```
tar xvf /dev/fd0
```

Where `/dev/fd0` is your floppy device.

4. Execute the following command to finish the installation:

```
/tmp/init.edge
```

You will see a window similar to Figure 234.




Figure 234. Start of installation dialog

The installation program guides you through the installation process. The windows are intuitive. During the installation process, you can also configure your backup device(s) and your scheduling schema for unattended operation. If information is needed during this process, you are asked to enter the appropriate data.

Now you are ready to use the product.

The actions *Resource Manager* and *Defining Devices* can be started by entering on the command line:

```
/usr/lib/edge/bin/edge.resmgr (Resource Manager) or  
/usr/bin/edge.config (Defining Devices)
```

You can also perform these actions, if you click **Admin** on BackupEDGE's main window.

16.2.2 Initializing the tape

Before you start making backups you should initialize the tape. To do this, you follow these steps:

1. Start the edgemenue program by executing command:

```
edgemenue
```

You will see a window similar to Figure 235.




Figure 235. BackupEdge main menu

2. In the Admin menu select **Initialize Tapes**. You will see a window similar to Figure 236.

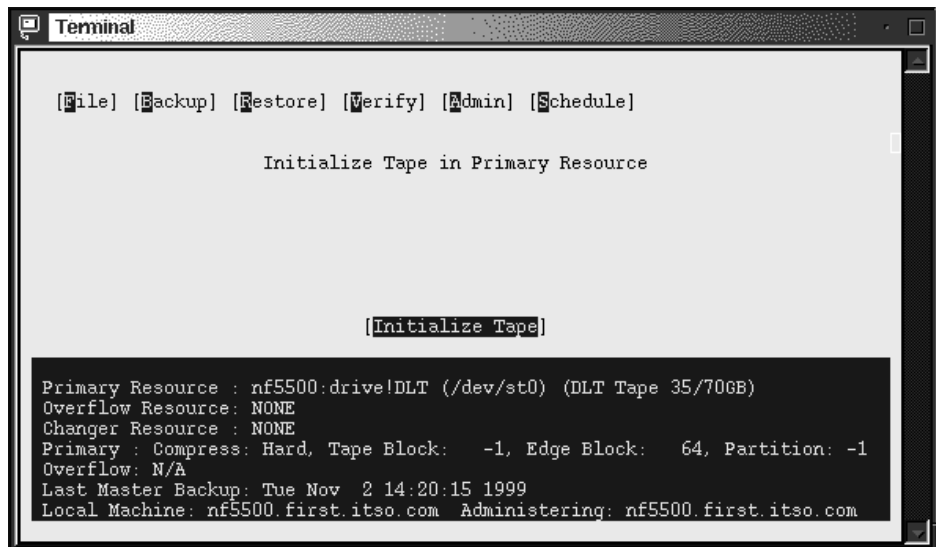


Figure 236. Initializing the tape

3. Select **Initialize Tape** and press Enter. The tape will be initialized. You will get a message that the tape is successfully initialized. Press Enter to continue.

You can check the tape properties by selecting **Show Tape Label** in the Verify menu. You will see a window similar to Figure 237.



Figure 237. Tape information

16.2.3 Your first backup

In this section we will show how to make backups of desired files or directories. You can perform backups in the edgemenu utility. Follow these steps to make a sample backup:

1. Start the edgemenu program by executing the following command:

```
/usr/bin/edgemenu
```

You will see a window similar to Figure 238.




Figure 238. Starting the backup

2. In the Backup menu select **Backup Files / Dirs**, and you will see a window similar to Figure 239.




Figure 239. Selecting source for backup

3. In the Files / Directories to Include field, type in the files or directories you want to back up. In our example we want to make backups of the directory /usr/src. Select **OK** to continue. You will see a window similar to Figure 240.




Figure 240. Backup in progress

After the backup is finished you will see a window similar to Figure 241.

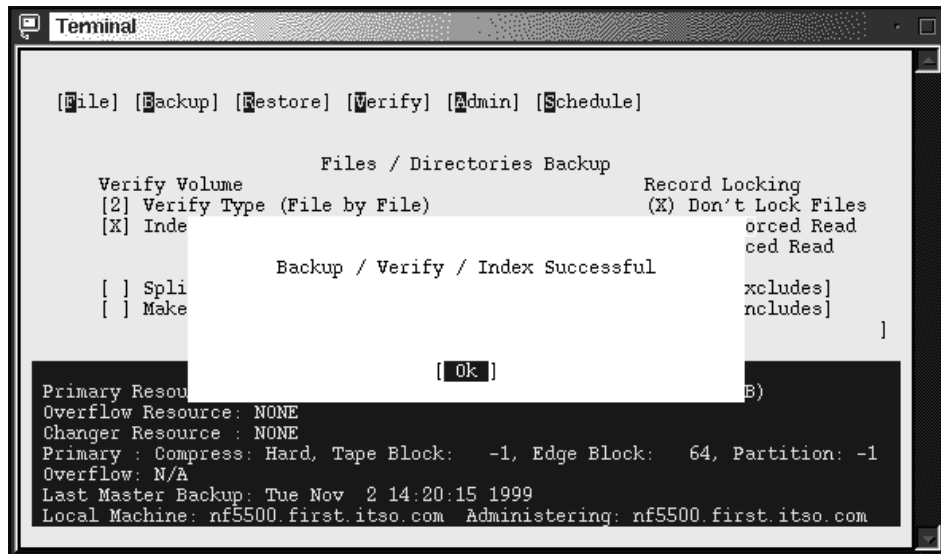


Figure 241. Backup completed

You will also see the backup report similar to Figure 242.




Figure 242. Backup report

You have just made your first backup and your files are safe now!

16.2.4 Restoring single files or directories

In this section we will show how to recover files from the backup. We are assuming that you are recovering files on the same server you made backups with the same user ID. You can perform recovery from the same utility as backups. Follow these steps to recover files:

1. Start the edgemenu program by executing the following command:

```
edgemenu
```

You will see a window similar to Figure 238. Select **Restore** and a window similar to Figure 243.




Figure 243. Starting the recovery

2. Select **Restore > Individual Files**, and you will see a window similar to Figure 239 on page 293.
3. Select the files or directories to restore. Select **OK** to continue, and you will see a window similar to Figure 244.




Figure 244. Recovery in progress

When the recovery is completed you will see a window similar to Figure 245.


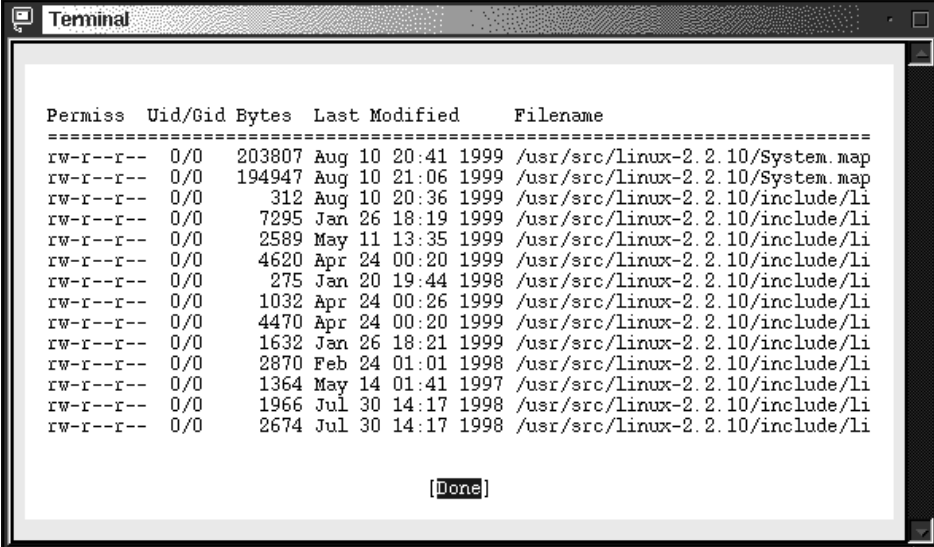


Figure 245. Recovery completed

Select **OK** to continue and you will see a recovery report similar to Figure 246.



```
Terminal
-----
Permiss Uid/Bid Bytes  Last Modified  Filename
-----
rw-r--r-- 0/0    203807 Aug 10 20:41 1999 /usr/src/linux-2.2.10/System.map
rw-r--r-- 0/0    194947 Aug 10 21:06 1999 /usr/src/linux-2.2.10/System.map
rw-r--r-- 0/0      312 Aug 10 20:36 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     7295 Jan 26 18:19 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     2589 May 11 13:35 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     4620 Apr 24 00:20 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0      275 Jan 20 19:44 1998 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     1032 Apr 24 00:26 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     4470 Apr 24 00:20 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     1632 Jan 26 18:21 1999 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     2870 Feb 24 01:01 1998 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     1364 May 14 01:41 1997 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     1966 Jul 30 14:17 1998 /usr/src/linux-2.2.10/include/li
rw-r--r-- 0/0     2674 Jul 30 14:17 1998 /usr/src/linux-2.2.10/include/li
-----
[Done]
```

Figure 246. Recovery report

Your files were recovered successfully!

16.2.5 Master and incremental backups

Usually system administrators perform so-called master and incremental backups. The master backup is a backup of all files on the system. Incremental backup is a backup of only those files that have changed from the last master backup. When you need to restore your data, restore the master backup and the last incremental backup. BackupEDGE can perform different types of incremental backups. Refer to the BackupEDGE manual for the explanation of them. Master and incremental backups can be performed from the edgemenu utility.

To perform a master backup follow these steps:

1. Start the edgemenu program by executing the following command:

```
edgemenu
```

You will see a window similar to Figure 238 on page 293.

2. Select **Backup > Master Backup**, and you will see a window similar to Figure 247.




Figure 247. Starting the master backup

3. Choose the options you want and select **Execute Backup** to start the backup. You will see a window similar to Figure 248.

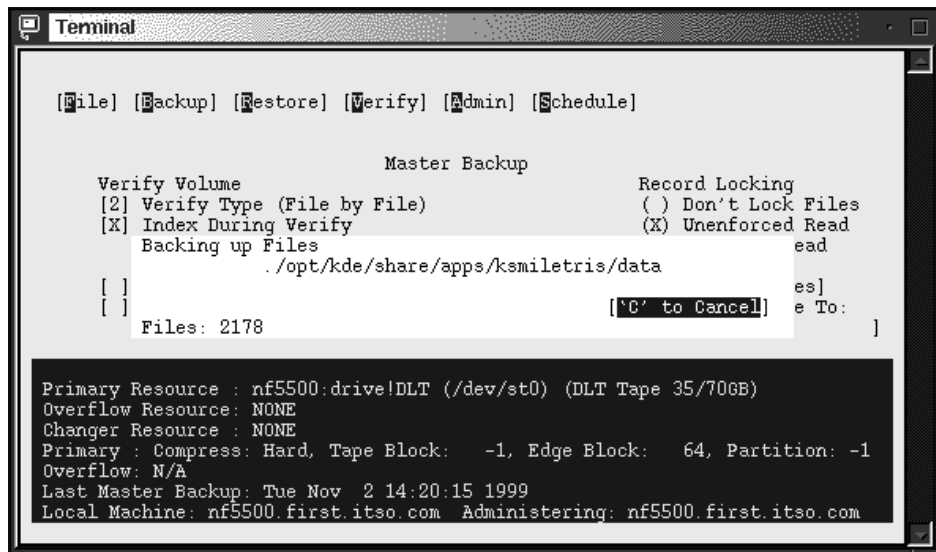


Figure 248. Master backup in progress

When the backup is finished you will see a window similar to Figure 249.

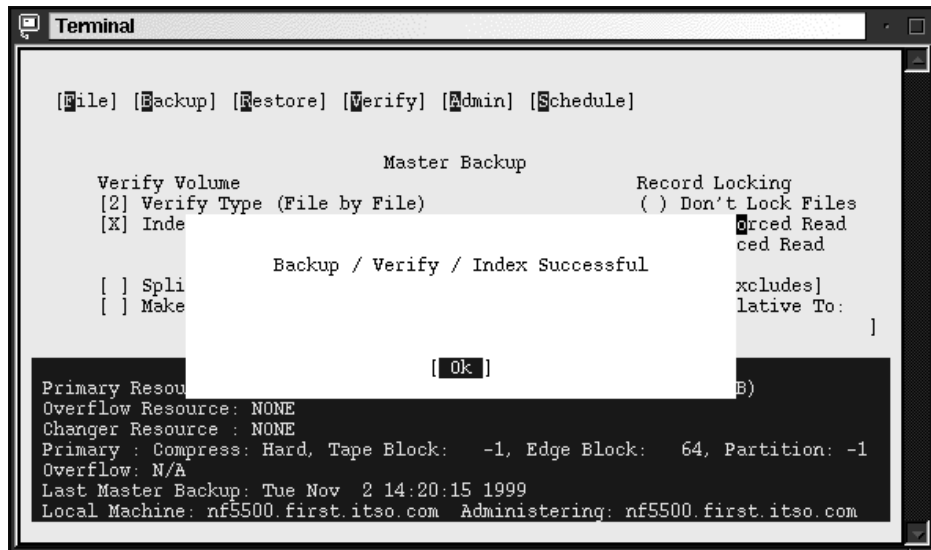


Figure 249. Master backup completed

Select **OK** to finish the operation, and you will see a backup report similar to Figure 250.

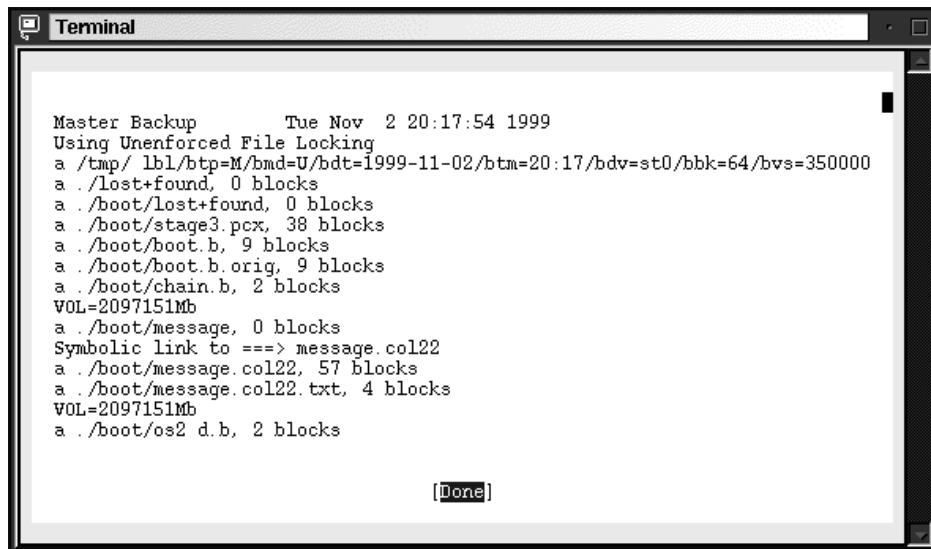


Figure 250. Master backup report

To perform incremental backups select **Backups > Incremental Backup**. Then follow the instructions in the window; they are similar to the ones for master backup.

16.2.6 Restoring master and incremental backups

To restore master and incremental backups you can use the edgemenue utility. When you start the utility and choose **Restore** you will see a window similar to Figure 251.

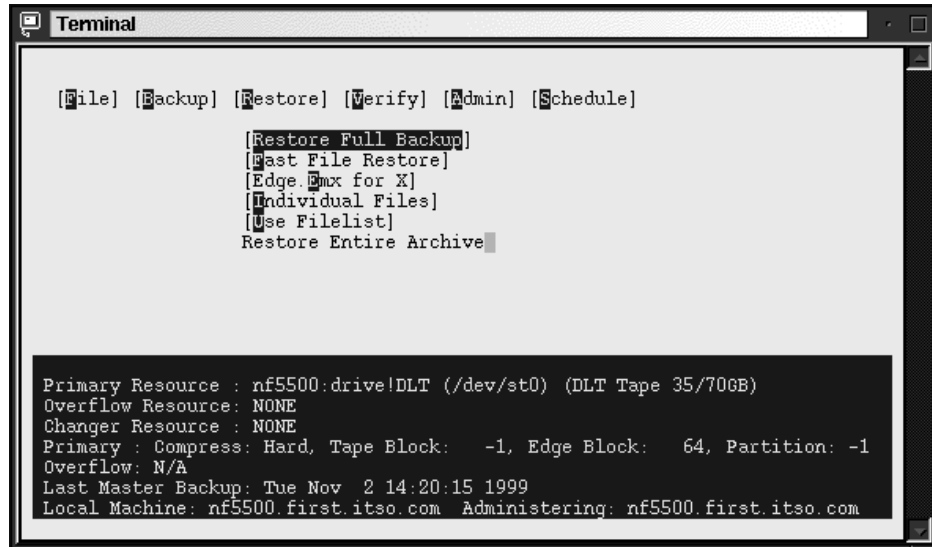


Figure 251. Starting restore full backup

Select **Restore > Restore Full Backup** and you will see a window similar to Figure 252.

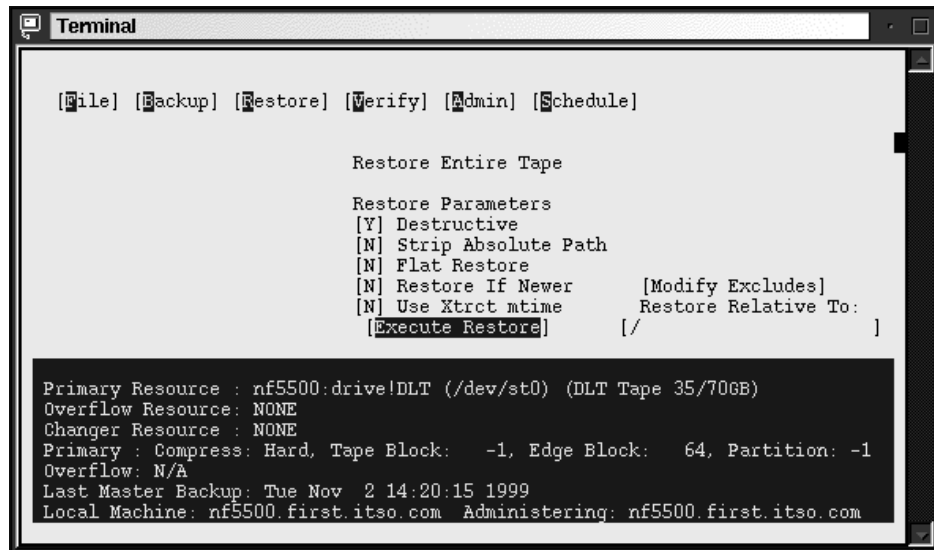


Figure 252. Full backup restore options

Choose your options and select **Execute Restore** to start restoring files.

16.2.7 Performing scheduled backups

To perform scheduled backups, you can use the `edge.nightly` utility included with BackupEDGE. To start this utility, execute the command:

```
/usr/lib/edge/bin/edge.nightly
```

But before you can use scheduled backups, you need to define them. To do this follow these steps:

1. Start the `edgemenu`.
2. Select **Schedule > Nightly Scheduling**. You will see a window similar to Figure 253.

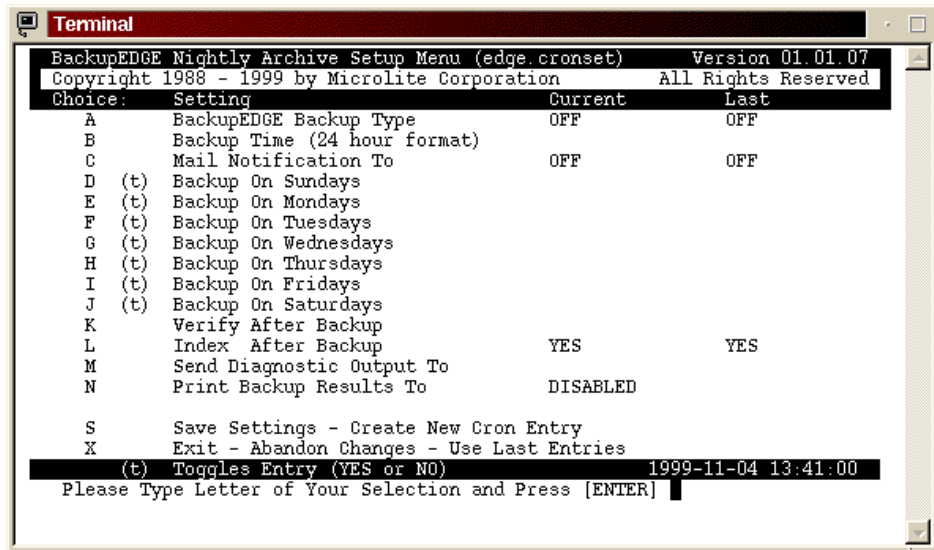


Figure 253. Schedule setup

- Here you can define the schedule for your backups. You need to define the type and time of the backup. To define the type of the backup select **A** and press Enter, and you will see a window similar to Figure 254.




Figure 254. Defining the type of backup

- Specify the type of backup you want to perform. In our example we selected **M** for master backup. You will be returned to the main window.

Note

You cannot mix master and incremental backups. If your master backup fits on one tape cartridge, we recommend that you do a master backup daily. If your master backup will not fit on one tape cartridge, do a manual master backup once a week and do incremental backups daily.

- Next you need to specify the time of everyday backup by selecting **B** and pressing Enter. You will see a window similar to Figure 255.




Figure 255. Setting the time

- Define the time for your backups. You will see a window similar to Figure 256.

```

Terminal <2>
BackupEDGE Nightly Archive Setup Menu (edge.cronset)          Version 01.01.07
Copyright 1988 - 1999 by Microlite Corporation              All Rights Reserved
Choice:  Setting                                           Current      Last
A       BackupEDGE Backup Type                            Master      Master
B       Backup Time (24 hour format)                       12:30      12:30
C       Mail Notification To                               root       root
D (t)   Backup On Sundays                                 YES        YES
E (t)   Backup On Mondays                                 YES        YES
F (t)   Backup On Tuesdays                              YES        YES
G (t)   Backup On Wednesdays                            YES        YES
H (t)   Backup On Thursdays                            YES        YES
I (t)   Backup On Fridays                                 YES        YES
J (t)   Backup On Saturdays                              YES        YES
K       Verify After Backup                               BIT        BIT
L       Index After Backup                                YES        YES
M       Send Diagnostic Output To                         /dev/null
N       Print Backup Results To                           DISABLED

S       Save Settings - Create New Cron Entry
X       Exit - Abandon Changes - Use Last Entries
(t)    Toggles Entry (YES or NO)                          1999-11-04 12:27:39
Please Type Letter of Your Selection and Press [ENTER]

```

Figure 256. After schedule definition

7. Select **S** and press Enter to save the settings. The configuration program will create an entry in the cron database for executing the `edge.nightly` utility. From now on, cron will execute the backup utility as you defined in the previous steps.

Note

Before you start using scheduled backups, check if you need to copy the file `/usr/lib/edge/bin/S88egde` to the `/etc/rc.d/rc2.d` directory. This script will clear all zombie PIDs from the `edge.nightly` on the system restart.

You can also start `edge.nightly` from your own scripts. When you start it from a command line or a script, you have to be logged in as root. After `edge.nightly` is started it will perform an immediate backup.

16.2.8 Configuring the tape devices

Any time after installation you can define or change your backup device. To accomplish this follow these steps:

1. Start the `edge.resmgr` resource manager by executing the command:

```
/usr/lib/edge/bin/edge.resmgr
```

You will see a window similar to Figure 257.

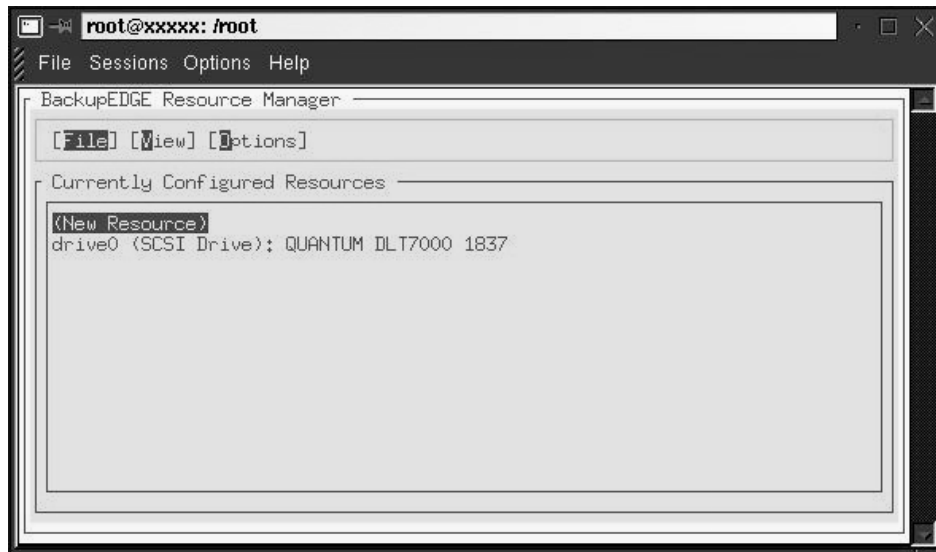


Figure 257. Starting the resource manager

2. Select **New Resource** and press Enter. You will see a window similar to Figure 258.

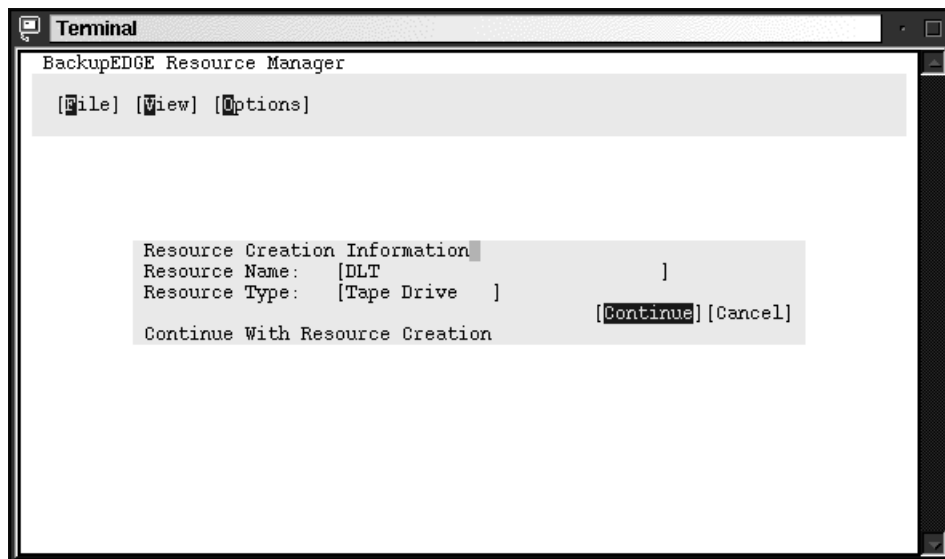


Figure 258. Defining the resource name

3. Type in the resource name and select a resource type. Select **Continue** to go on. You will see a window similar to Figure 259.

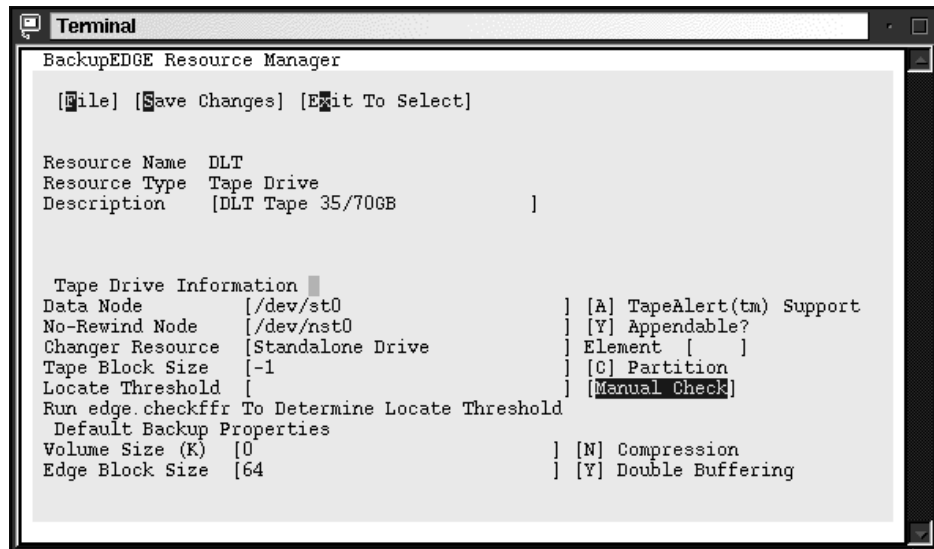


Figure 259. Parameters for the tape

4. Type in the description, data node and no-rewind node. In our example, the data node is `/dev/st0` and no-rewind node is `/dev/nst0`. You can leave all other fields as default.
5. Select **Manual Check** to define other parameters automatically. You will see a window similar to Figure 260.

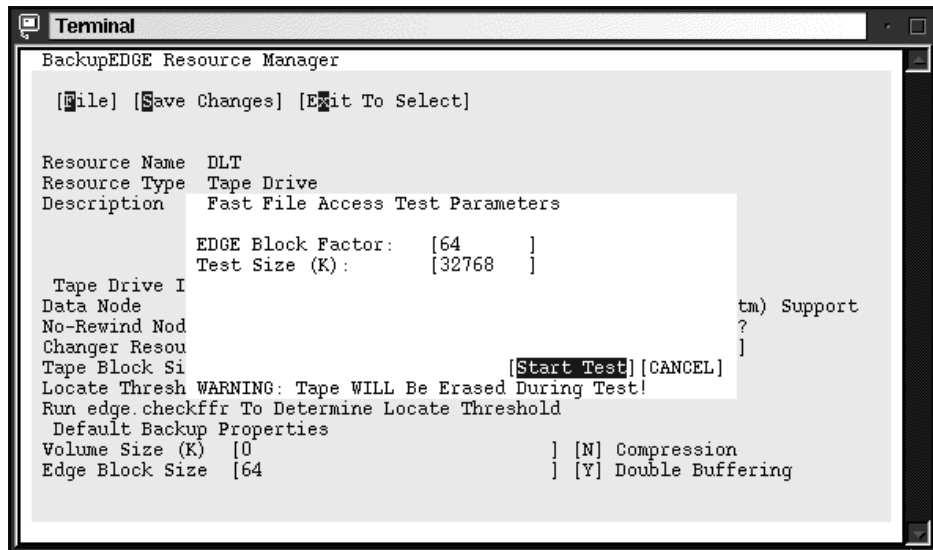


Figure 260. Setting the parameters for tests

- Here you can select the block factor and the test size. Select **Start Test** to continue. You will see a window similar to Figure 261.

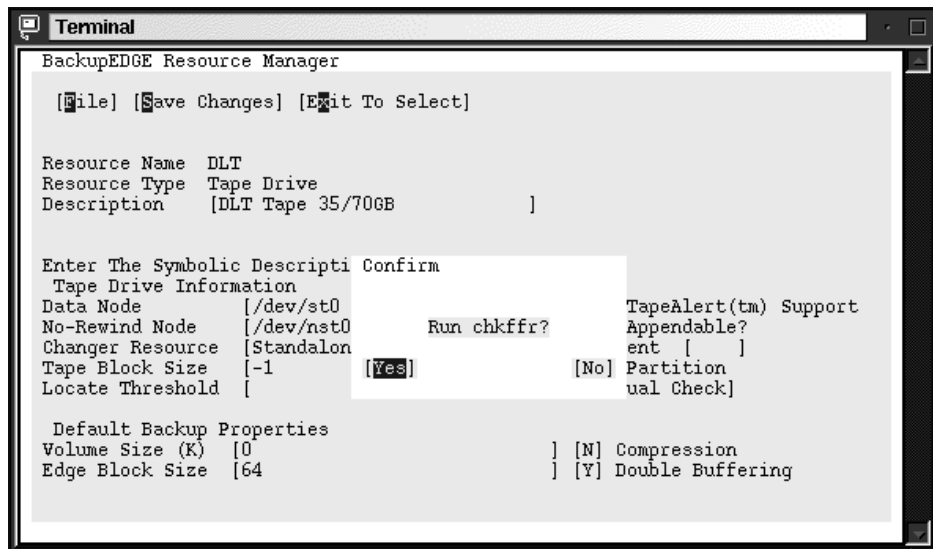


Figure 261. Starting the test

Stop

Performing this test will destroy all data on the tape.

7. Select **Yes** to continue. You will see a window similar to Figure 262.

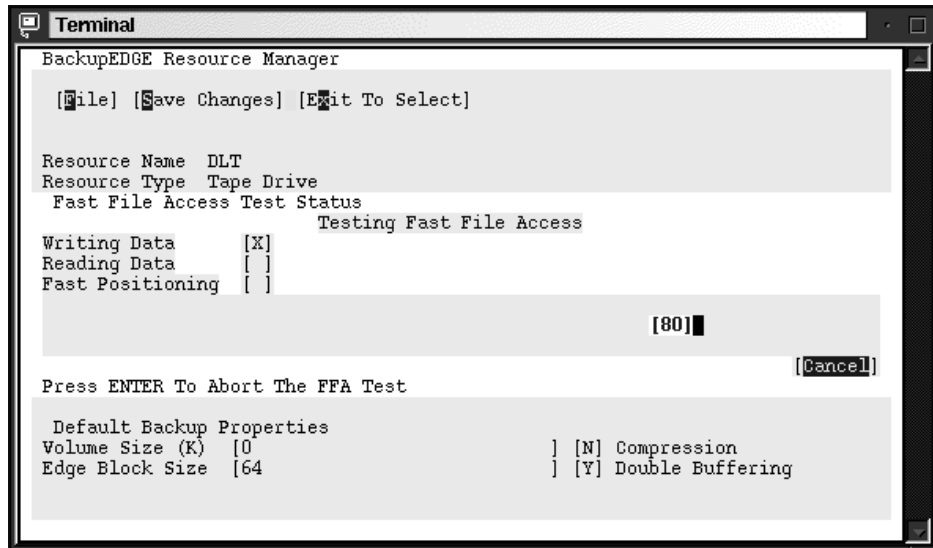


Figure 262. Performance test

After the test is done you will see a window similar to Figure 263.

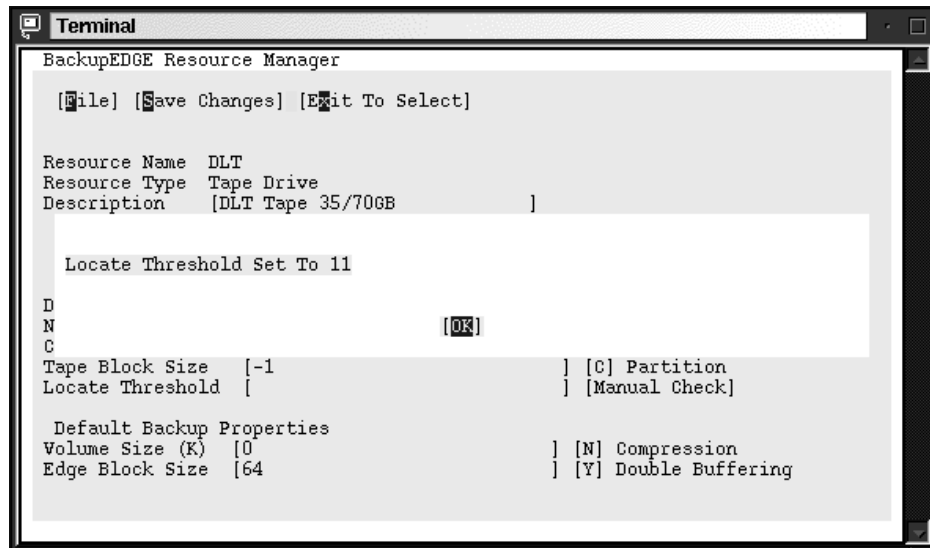


Figure 263. Threshold value

8. After the test is done you will see the proposed value for the threshold. Click **OK** to continue. You will be back in the parameters definition window similar to Figure 260 on page 308. Here you need to define four more parameters:
 - Volume Size
 - EDGE Block Size - the default size is 64 for a 32 KB buffer
 - Compression
 - Double Buffering - with multiple buffers you can increase the backup speed
9. Save the changes by selecting **Save Changes**. You will see a window similar to Figure 264.




Figure 264. Saving the device definitions

16.2.9 Defining the devices for making backups

Any time after installation when you configured your backup hardware device, you can change which device the backup software uses for each user performing backups. If you are logged in as root, you will define devices for the root user. Usually this is the only user doing backups on the system. Follow these steps to enter the resource manager for backup:

1. Start the edge.config configuration menu by executing the command:

```
/usr/bin/edge.config
```

You will see a window similar to Figure 265.

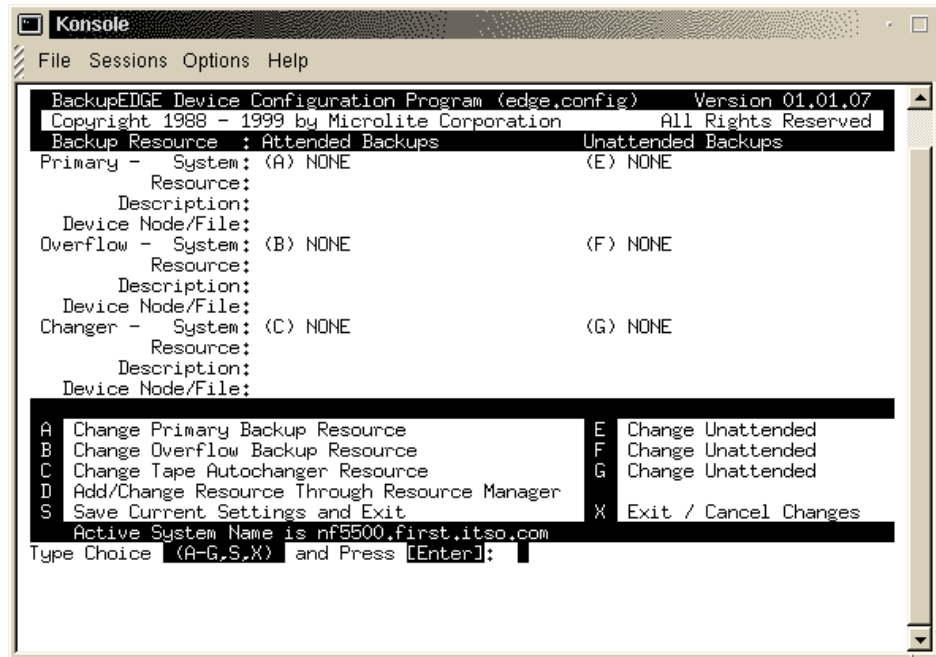


Figure 265. Device Configuration

2. Here you need to define the devices for attended and unattended backups.
3. Type in A and press Enter to define the device for attended backups. You will see a window similar to Figure 266.

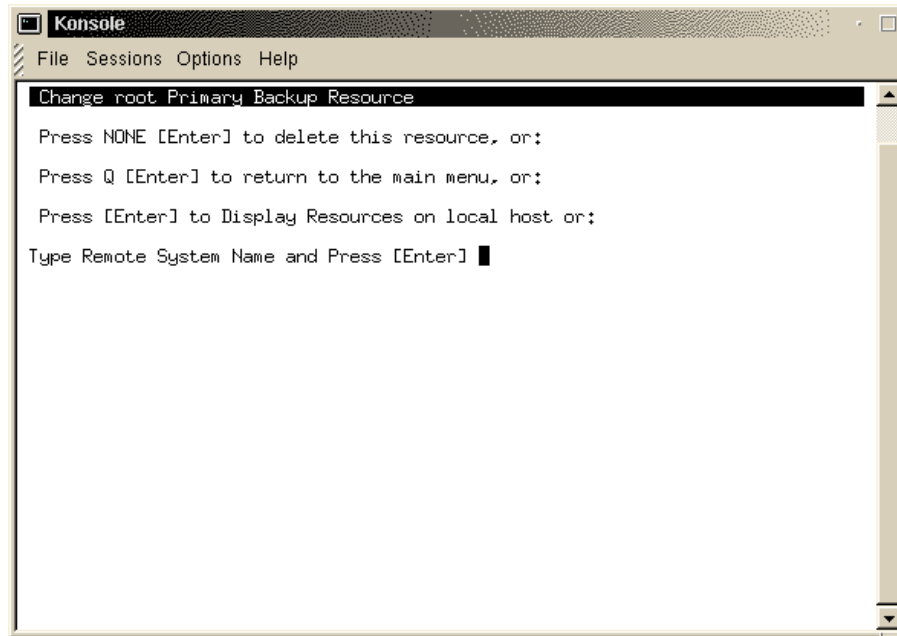


Figure 266. Selecting the device for backup

4. Press Enter to continue. In the next window you will see all defined backup devices. Type in the device you want and press Enter to continue. You will see a window similar to Figure 267.

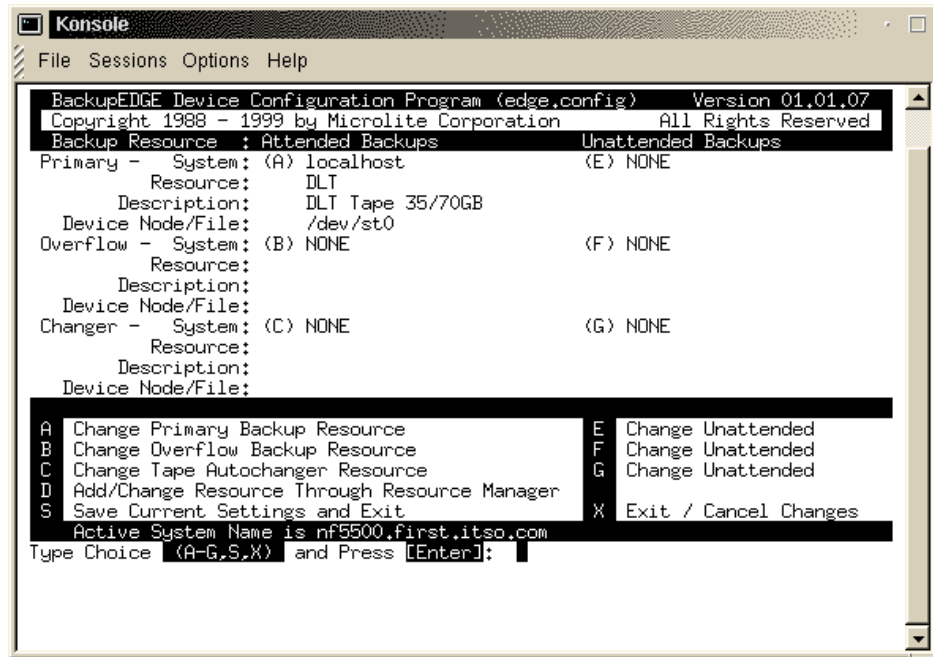


Figure 267. After definition of attended backup device

5. Follow the steps from 1- 4 for the unattended device also.

16.2.10 Microlite RecoverEDGE

By using the RecoverEDGE tools you can create emergency recovery diskettes to rebuild your system in the case of disaster. RecoverEDGE handles the details of reconstructing your FDisk, divvy, and/or slice tables, rebuilding your file systems and restoring your data, even if your hard drive size has changed. RecoverEDGE uses your live system backups, so there is no need to shut down your system in order to protect it. You can even restore your system over the network.

With RecoverEDGE restoring the system is very easy. To recover the system you should follow these tasks:

1. Identify and correct the cause of the failure.
2. Boot from the RecoverEDGE disks.
3. Reconfigure your file systems.
4. Restore your backups.
5. Shut down and reboot.

6. System is ready to use.

Note

RestoreEDGE uses your master and incremental backups for recovery, so the accuracy of the data depends on these backups.

16.2.10.1 Creating the RecoverEDGE boot disks

Before you can use RecoverEDGE for disaster recovery you should build a set of boot disks. To create the boot disks follow these steps:

1. Start the utility for creating the RecoverEDGE boot diskettes:

```
/usr/bin/re2
```

or go to **Admin>Make RecoverEDGE Media** in the menu.

You will see a window similar to Figure 268.

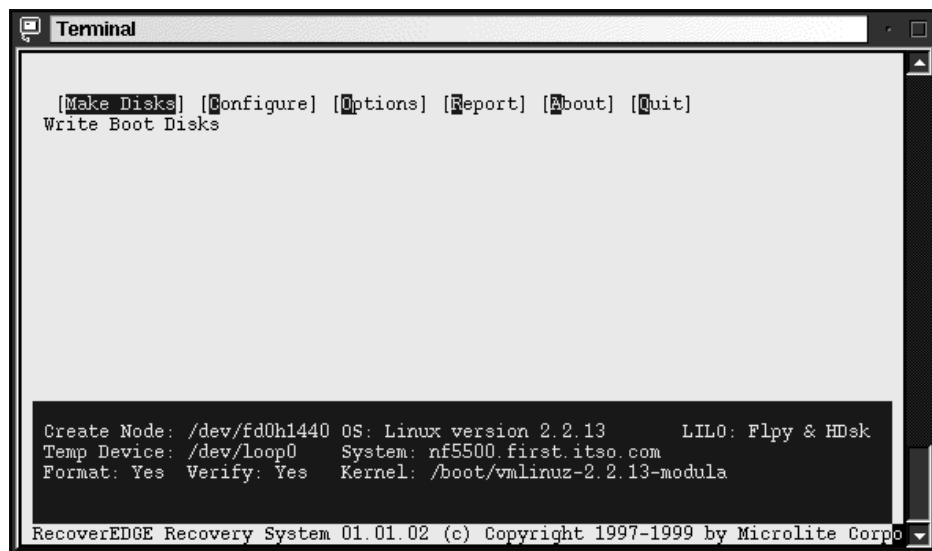


Figure 268. RecoverEDGE utility

2. Select the **Configure** option and press Enter, and you will see a window similar to Figure 269.




Figure 269. Configure menu

3. Select the **Disk Layout** option and press Enter, and you will see a window similar to Figure 270.




Figure 270. Disk layout menu

- Here you can configure the kernel, modules, network and the file systems for your RecoverEDGE boot disks. Select the **Kernel** option and you will see a window similar to Figure 271.

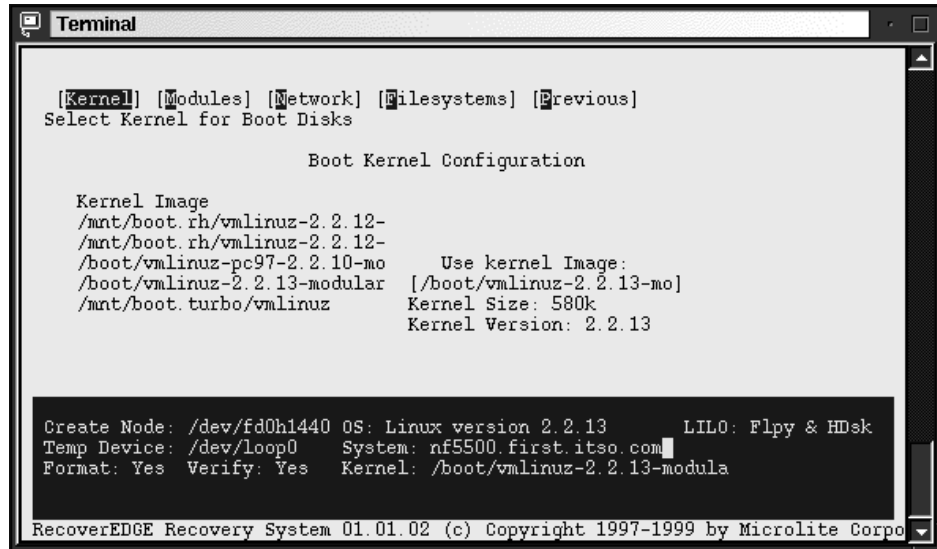


Figure 271. Kernel options

Here you define which kernel will be used for creating the diskette.

- Return to the previous stage and select **Modules** and press Enter, and you will see a window similar to Figure 272.

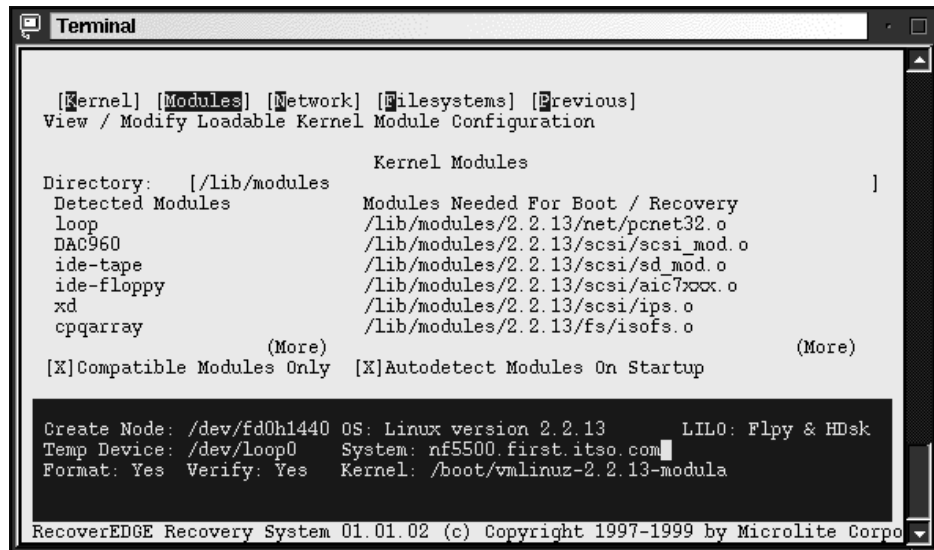


Figure 272. Modules options

Here you define which modules will be used for building the initial RAM disk for the recovery system. In the Directory field you can specify the path to the modules that corresponds to the kernel you defined for booting. If you choose the option **Autodetect Modules on Startup**, RecoverEDGE will load currently loaded modules.

Note

Do not forget to include the module for the tape drives.

- Return to the previous stage and select **Network** and press Enter, and you will see a window similar to Figure 273.

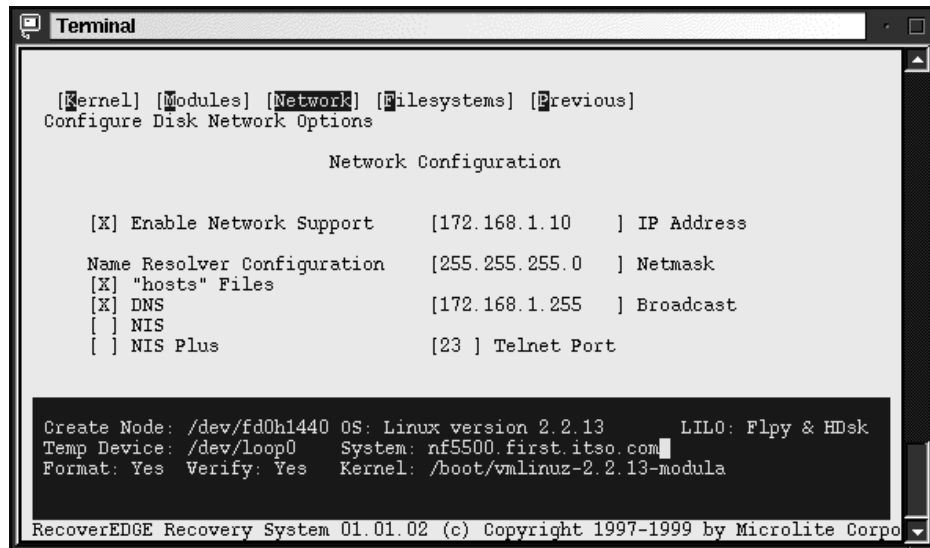


Figure 273. Network options

Here you define your network setup in case you will restore the system from a tape device on the network. You do not need this if you have a locally attached tape.

7. Return to the previous stage and select **Filesystems** and press Enter, and you will see a window similar to Figure 274.




Figure 274. Filesystems options

Here you define which mounted file systems will be recovered.

8. Return to the configuration panel and select the **Boot Loader** option and press Enter. You will see a window similar to Figure 275.

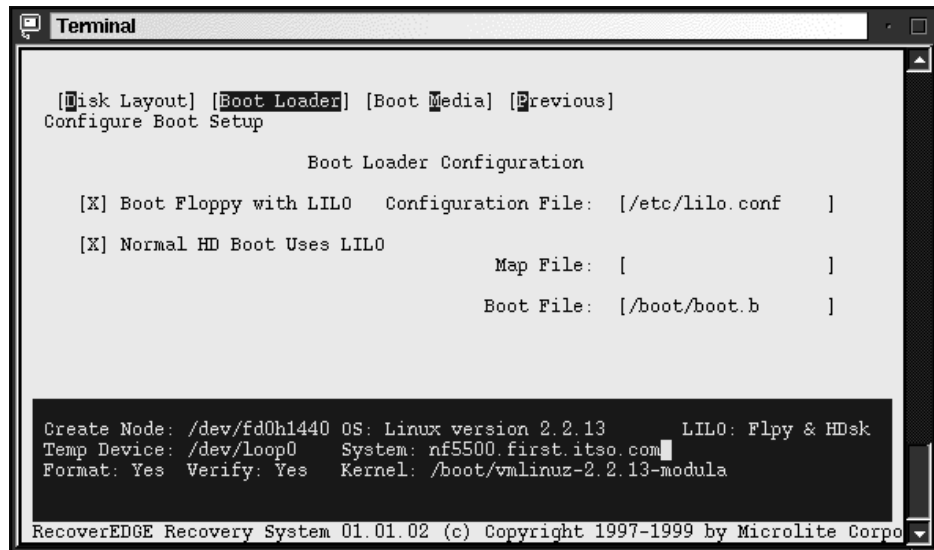


Figure 275. Boot Loader options

Here you define options for the Boot Loader.

- Return to the configuration panel and select the **Boot Media** option and press Enter. You will see a window similar to Figure 276.




Figure 276. Boot Media options

Here you define how the boot diskettes will be created.

10. After you configured all settings return to the main window and select **Make Disks**. You will be prompted to insert three diskettes.

Note

If you get an error that diskettes cannot be created, the probable cause is that images are too big. Try to reduce the number of loaded modules or even make the special kernel just for this purpose, throwing out all unnecessary things.

After the diskettes are created you are ready to deal with disaster on your system. But before this really happens, try to boot from these diskettes and verify if your tape device is recognized.

16.2.10.2 Verifying the RecoverEDGE boot diskettes

To verify the diskettes, boot from the first diskette and follow instructions on the window. When the system is started you will get the RecoverEDGE main menu. Select **Utilities > Tape Drive**.

In the Tape Device Node field, you see the defined tape device. Go to the Test Tape Drive field and test your tape device. If the test is successful your recovery set is ready to use.

16.2.10.3 Recovering from a total crash

To recover from a disaster crash follow these steps:

1. Resolve all hardware problems.

Note

Before restoring the system, initialize the Master Boot Records of all disk drives.

2. Boot the server from the first RecoverEDGE boot diskette.
3. When you are prompted to insert the root diskette, insert the second RecoverEDGE boot diskette. After the diskette is loaded, RecoverEDGE will start and you will see a window similar to Figure 277.




Figure 277. RecoverEDGE initial window

4. Select **Restore > One Touch**. Follow the instructions on the window to complete the recovery.



Note

For recovery you will use your master and incremental backups.

5. When all files are backed up, press a key to get back to the main window. All the file systems will be then synchronized and LILO will be set up and executed.
6. Before you reboot, switch to a console 2 with Alt+F2 and execute the following commands to check the fstab file for correct entries for your system:

```
mount /dev/sdb6 /mount  
cat /mount/etc/fstab
```

In our example `sdb6` is our root partition. You should use your root partition here.

That is all there is to it. Your restored system is ready to use.

16.2.11 More information on Microlite

For information on advanced features consult the *Microlite User's Guide* or the Microlite Web site at:

<http://www.microlite.com>

16.3 Arkeia

Arkeia is a complete client/server backup solution for Linux and other platforms. With Arkeia you can safely archive every file, directory, device node and special file on your file systems. Unlike standard UNIX tar command, which ignores many important files, Arkeia also verifies the data written to tape to ensure that the tape is an accurate reflection of your data. Below are the features provided by Arkeia backup software:

- Data Compression - automatic data compression is supported.
- GUI Interface. A CLI-interface is also available.
- The backup server may be your local system or a remote system.
- High Performance - advanced double buffering and variable block factors.
- Virtual File Support - you can back up virtual (sparse) files.
- Multi-Volume / Multi-Device Archives - automatic spanning across multiple volumes or devices.
- Wildcard Support - when selecting files you can use a wildcard.
- Raw Device Backups - you can archive an entire raw device/partition to tape.
- Master / Incremental Backups
- Unattended Operation - you can configure schemas to periodically perform full backups and/or incremental backups.

Arkeia is designed to operate on Linux kernels 2.x and there are available versions for several types of libraries (libc5 and libc6) and distributions.

Requirements for the server:

- A 486 processor or higher
- 32 MB RAM
- 1 GB disk space
- SCSI adapter card
- SCSI tape drive
- TCP/IP services
- Linux 2.0 or higher

Requirements for the client:

- A 486 processor or higher
- 5 MB disk space

In the following sections we describe how to install, configure and use the Arkeia backup software.

16.3.1 Installing Arkeia

Arkeia is available in different package formats (tar, rpm) for different distributions either on CD or downloadable from Arkeia's Web site (follow the link <http://www.arkeia.com>) in the DOWNLOAD AREA. To install Arkeia, we recommend that you follow the installation procedure described in the *Installation and Quick Start Manual*. You can find this manual on the Arkeia-CD or download it from Arkeia's Web site.

On the Arkeia server, you must also install the client and the GUI package. These packages are required to configure the backup server. After the installation of the client and GUI packages, you can install the server package.

16.3.2 Configuring Arkeia

Before you can configure Arkeia, check whether the Arkeia backup server is running. To do this, enter:

```
ps -ef | grep -v grep | grep nlservd
```

on the system which should be used as your backup server. If you see a line like

```
root 488 1 0 09:06 ? 00:00:00 /usr/knox/bin/nlservd start
```

the backup server is running. To begin with the configuration of Arkeia, be sure, you have X-Windows running. Then enter on the command line:

```
Arkeia
```

You will see a dialog like Figure 278:



Figure 278. Arkeia initial window

The field for the server name is by default filled in with the name of the system you currently work with. You must change this field if you have installed the server component on another system.

The field for the login name is by default filled in with `root`. Change it if you have changed the name of the Arkeia administrator.

The field for the password is empty by default. You have to enter the password when you have changed the password. The main dialog window of Arkeia appears (Figure 279):

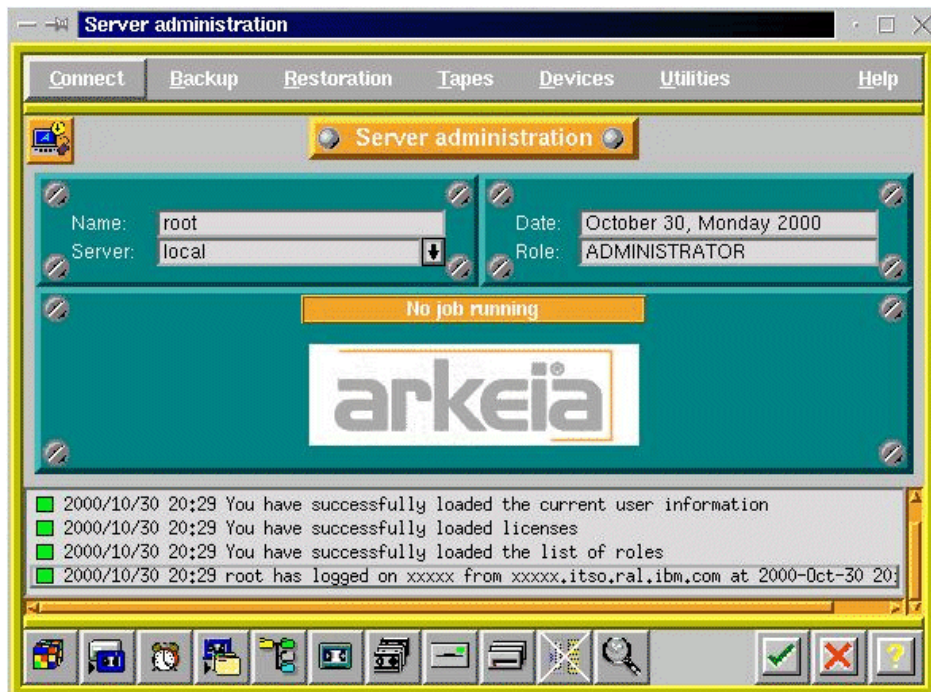


Figure 279. The Arkeia main dialog window

If you want a simpler layout of the window, go to **Utilities -> Setting** in the menu bar and modify the appearance of the windows. Click the **OK** button, save the new setting, and click the **OK** button again. Now, you will get a window similar to the window in Figure 280.

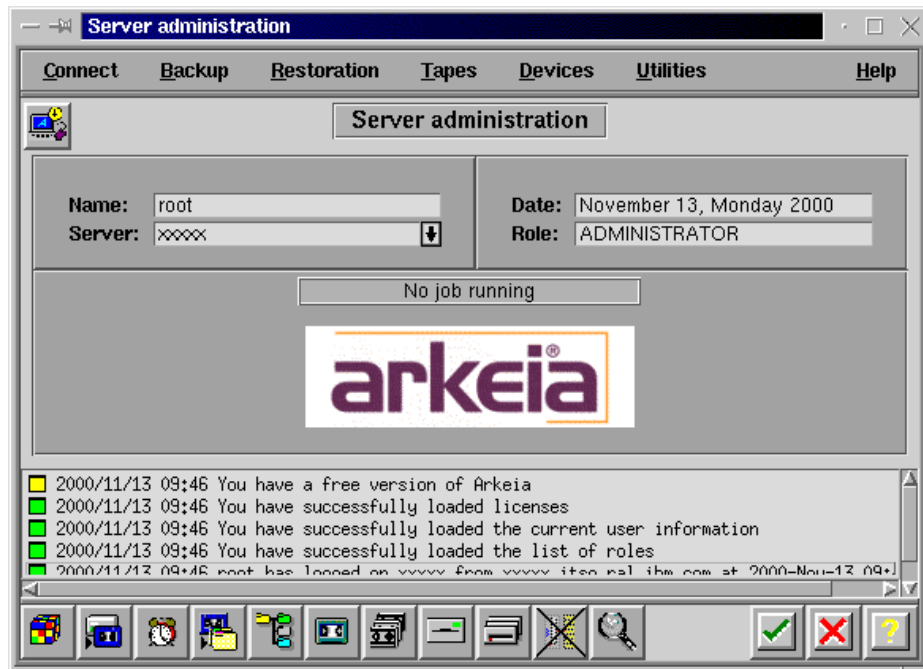


Figure 280. The new Arkeia main dialog window

At the bottom of the window you see push buttons shown in Figure 281:



Figure 281. Bottom part of main window

The meaning of these buttons is, from left to right:

- Refresh job
- Interactive backup
- Periodic backup
- Restoration
- Savepacks
- Tapes management
- Pools management
- Drives management
- Drivepacks
- Libraries management
- Backup done

- OK button. Clicking this button opens a new Welcome dialog.
- Cancel button. Clicking this button to leave Arkeia.
- Help

Before you can begin with your first backup, you must carry out the following configuration steps:

- Pool management
- Tape management
- Drives management
- Drivepacks management
- Savepacks management

Let us start with tape pool management. Click the pools management button on the bottom of the main dialog or click **Tapes -> Pools management** on the menu. The pools management window appears as in Figure 282:




Figure 282. Pools management main dialog window

The scratch pool exists by default. To create a new tape pool, for instance for your backup tapes, click the **new** button. The pool creation dialog appears as in Figure 283:




Figure 283. Pool creation window

Fill in the dialog fields with the appropriate information and click the **OK** button. The pools management main windows appears with the pool list updated as in Figure 284:




Figure 284. Pools management main window with updated pool list

To return to the main dialog, click the **OK** button. Now we can fill the Full Backup pool with tapes. To do this, click the tape management button or click **Tapes -> Tapes management** in the menu. The tapes management main window appears (Figure 285):

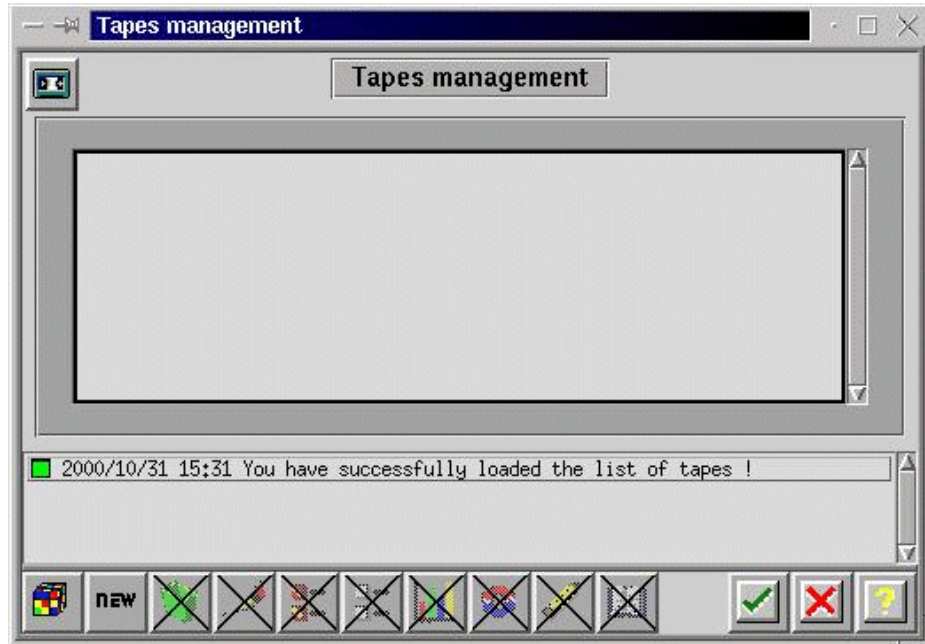


Figure 285. Tape management main window

Click the **new** button to enter new tapes (Figure 286):

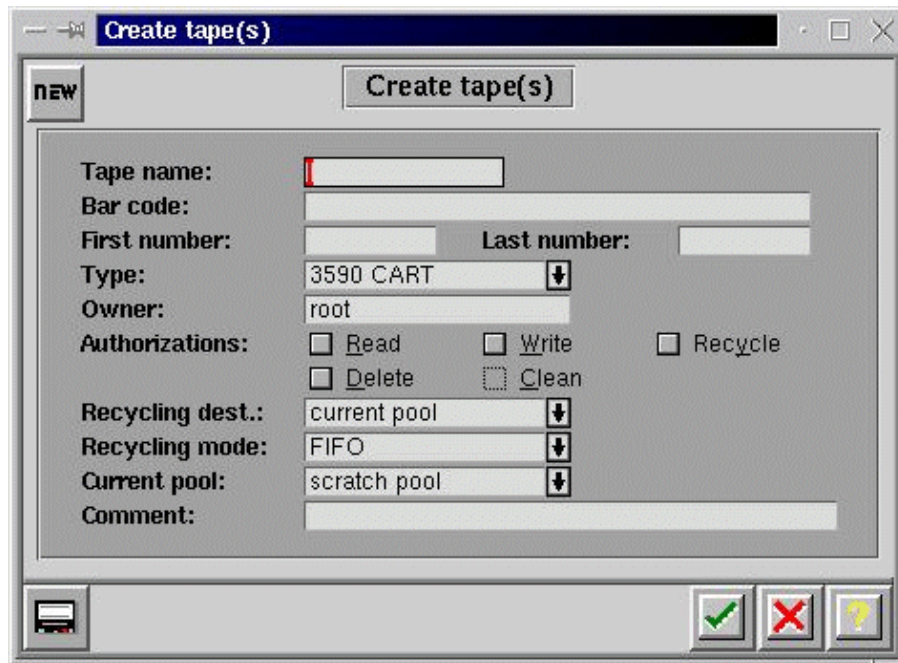


Figure 286. Create tape(s) window

The tape name consists of a fixed part and a variable part. The fixed part can be any text, while the variable part is a number. Enter the first part of the tape name, the first and the last number of the tapes to be used, and the tape type (DAT, DLT, etc.). Choose the pool these tapes should belong to and enter a comment in the comment line. Click the **OK** button to return to the tapes management main window. The tapes management main window appears with the updated list of currently created tapes. Click the **OK** button in this window to return to the main window.

After the creation of tape pools and tapes, we can create drives and drive packs.

Drives must be created first. To do this, click the drives management button in the main window or click **Devices -> Drives management** in the menu. The drives management window appears (Figure 287):




Figure 287. Drive management window

Click the **new** button to fill in the fields with the appropriate information. The fields Name and Rewind Device must be filled. Do not forget to choose the correct tape type in the Type field. To return to the Arkeia main window, double-click the **OK** button.

Now we can generate drivepacks. Press the drivepacks button or click **Devices -> Drivepacks** on the menu. The Drivepacks window appears (Figure 288):

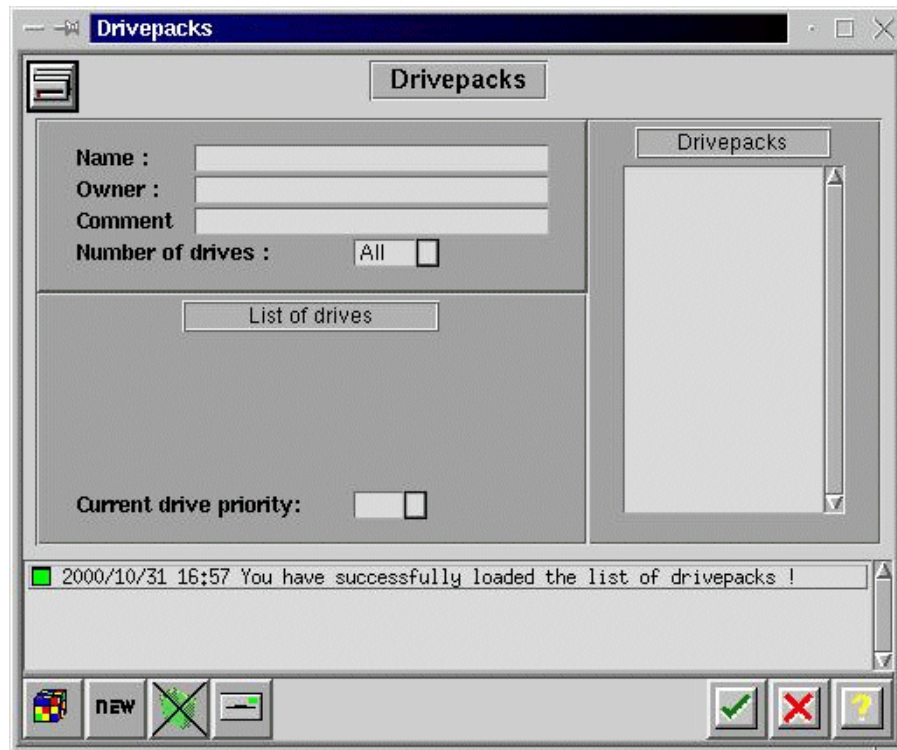


Figure 288. Drivepacks management window

Click the **new** button to fill in the fields. Fill in the Name field and choose one entry in the drives list and click the **OK** button to update the list of existing drivepacks on the right side of the window (Figure 289).

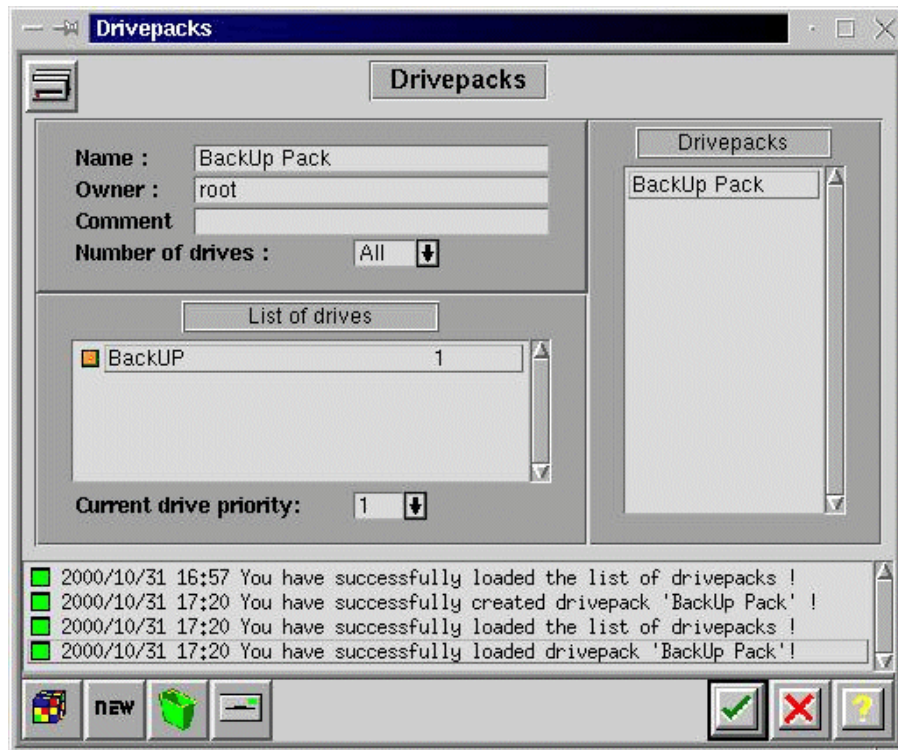


Figure 289. Updated drivepacks management window

Click the **OK** button again to return to the main dialog window.

The last step to be done before data can be saved is creating at least one savepack. You describe in savepacks which data should be saved. Different savepacks contain different sets of data to be saved.

To create savepack(s), click the Savepacks button or click **Tapes -> Savepacks** on the menu. You will see a window like Figure 290:

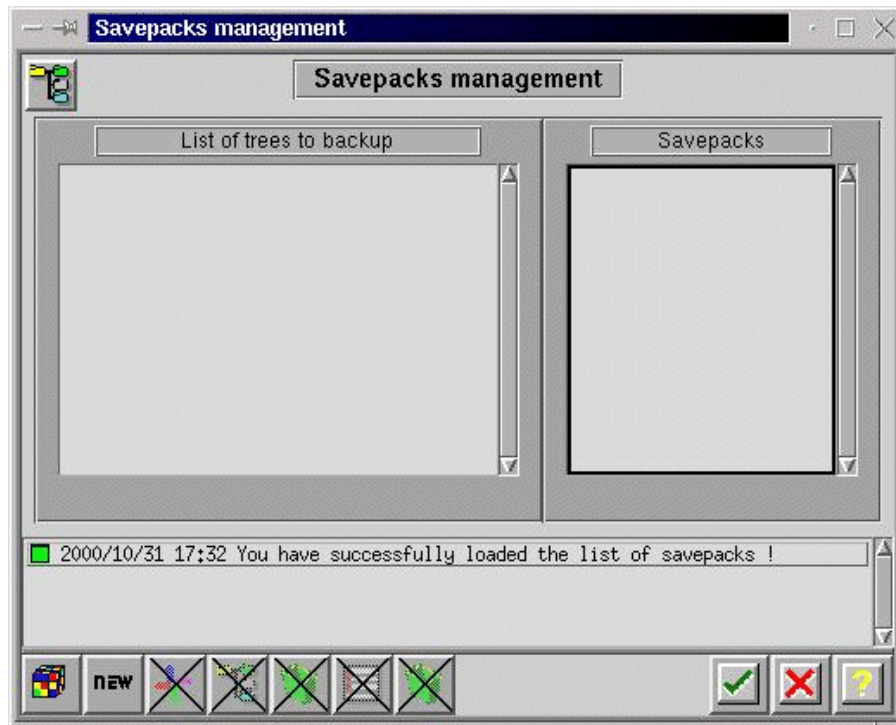


Figure 290. Savepacks management window

Click the **new** button to enter input mode. A window similar to Figure 291 appears.

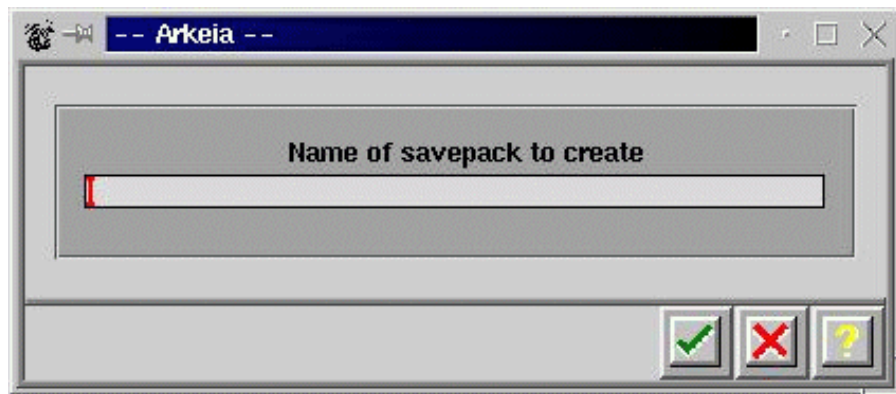


Figure 291. Window to create a new savepack

Enter the name of the new savepack and click the **OK** button to return to the updated savepacks management window (see the list of savepacks on the right side of the window). A window like in Figure 292 appears:

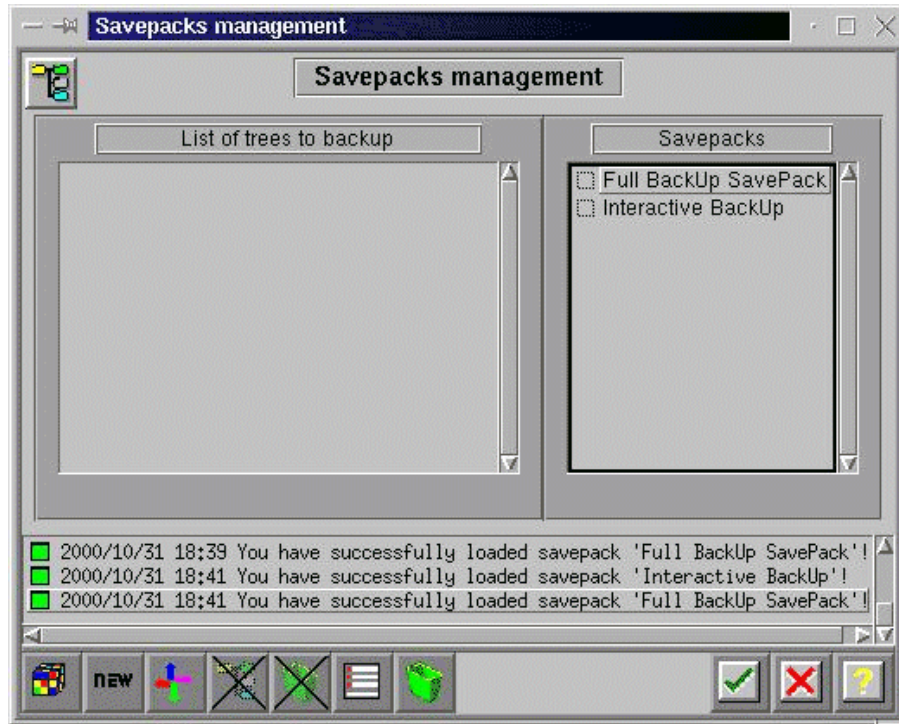


Figure 292. Updated savepacks management window

Now, you select the data that should be saved in every created savepack. Move the cursor over the name of the savepack you want to select the data for and click the left mouse button. You can see the selected savepack.

Now, move the cursor over the list of trees to back up (left listbox of this window), click the right mouse button and select **Navigator** in the upcoming pull-down menu. You will see a window similar to Figure 293.

To navigate through the directory tree of a system shown in this window, move the cursor over the system you want to select and double-click the left mouse button. A window similar to Figure 294 appears.

Double-clicking the left mouse button over a directory symbol opens this directory and shows the content of this directory.

Clicking once with the left mouse button in the checkbox to the left of a directory name or file name toggles the select/unselect status of this item. All selected items will be inserted in the list of trees to back up for the selected savepack. If you select a directory, the checkbox changes the color totally. If you select only a selection of the items in a directory, the checkbox for this directory changes color only in the right half of the checkbox.

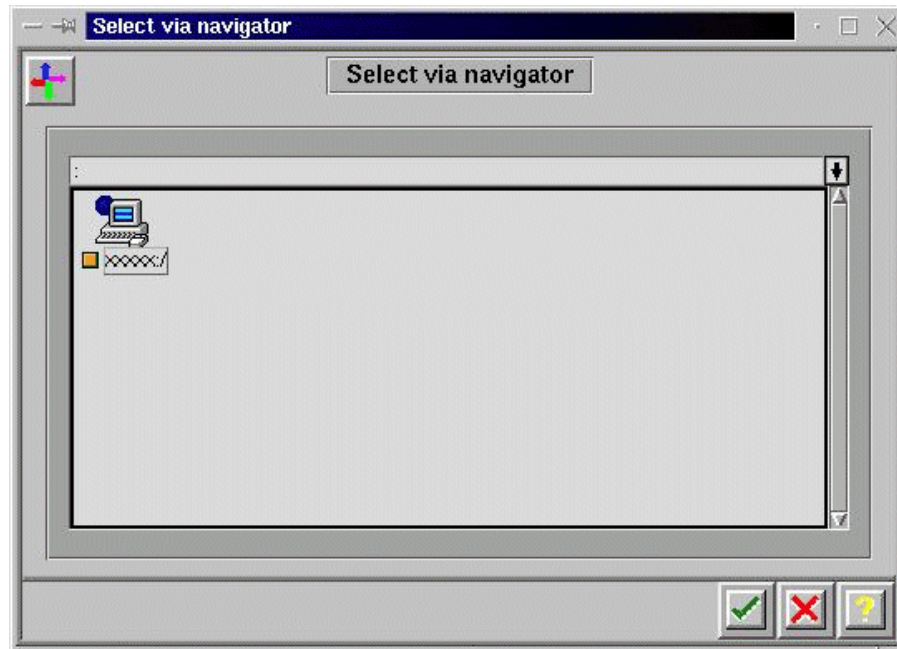


Figure 293. Navigator window



Figure 294. Updated navigator window

To return to the savepacks management window, click the **OK** button. You will see a window similar to Figure 295.

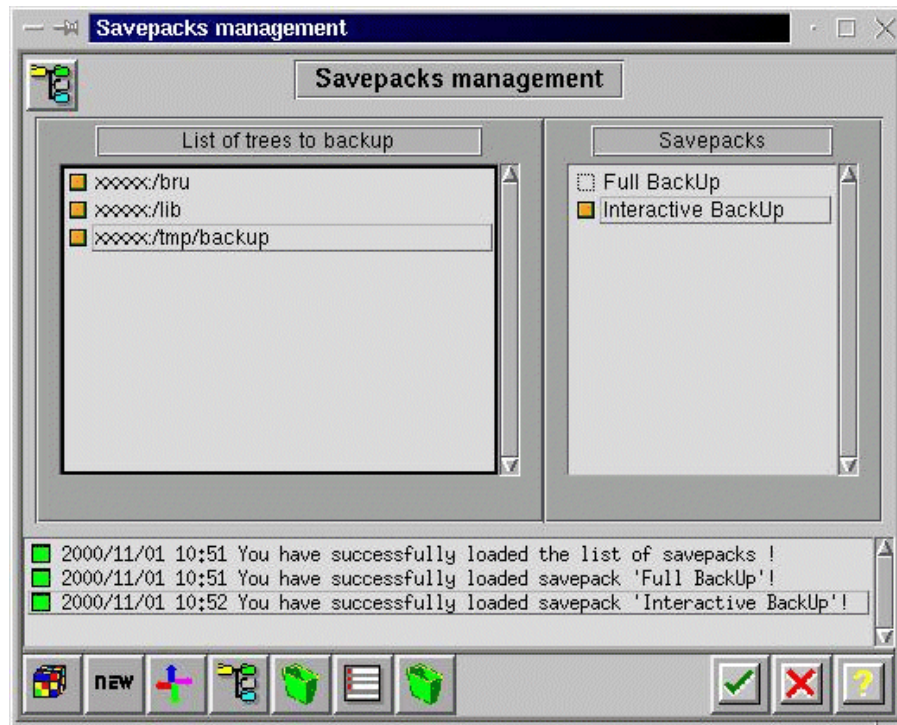


Figure 295. Updated savepacks management window

The basic configuration steps are now done.

Read the *Administrator's Manual* to get more information about the advanced possibilities of Arkeia.

16.3.3 Interactive backup

To start an interactive backup, click the interactive backup button or click **Backup>Interactive Backup** on the menu. A dialog like Figure 296 appears.




Figure 296. Interactive bckup start window

In the comboboxes Savepack, Drivepack and Pool fields, choose which data sets should be backed up on which tapes and on which tape drives.

In the **Type** box, choose between **Total Backup** and **Incremental Backup** and between **Standard** and **Continous**.

In the Tape Strategy field, choose between **Use new tapes** and **Complete existing tapes**.

In the **Valid for** field, decide how long the tape(s) for this backup should be valid.

Click the **OK** button to proceed. A window as in Figure 297 appears.

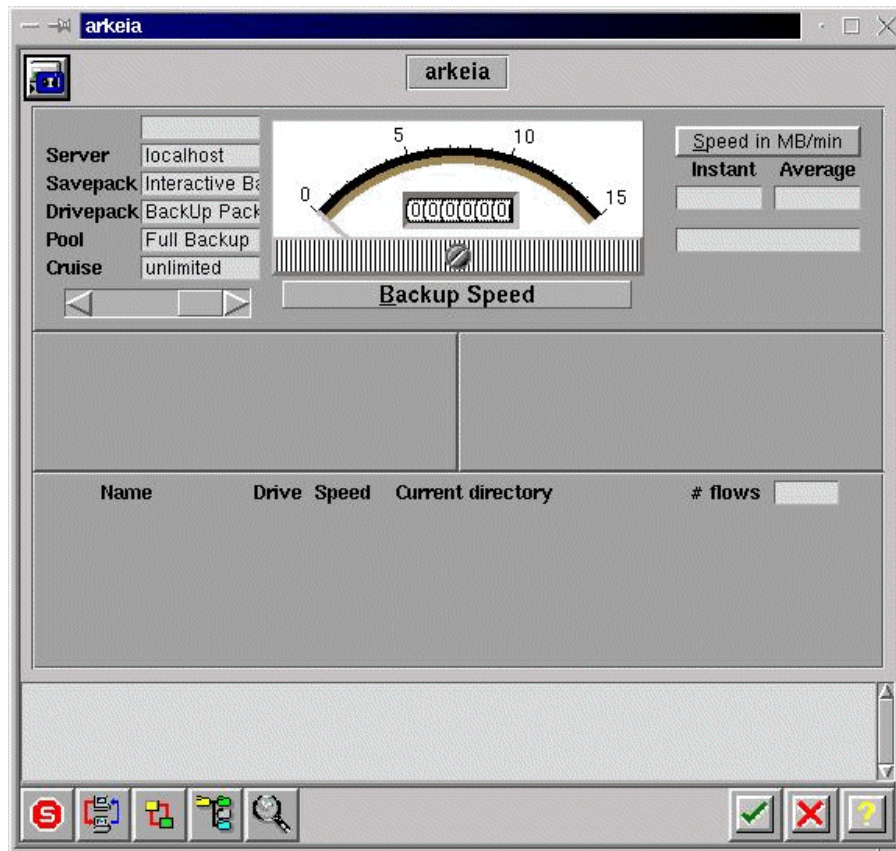


Figure 297. Arkeia's main window during backup

As the backup process proceeds, the content of this window will change. Most of the time, you will see a window like Figure 298.

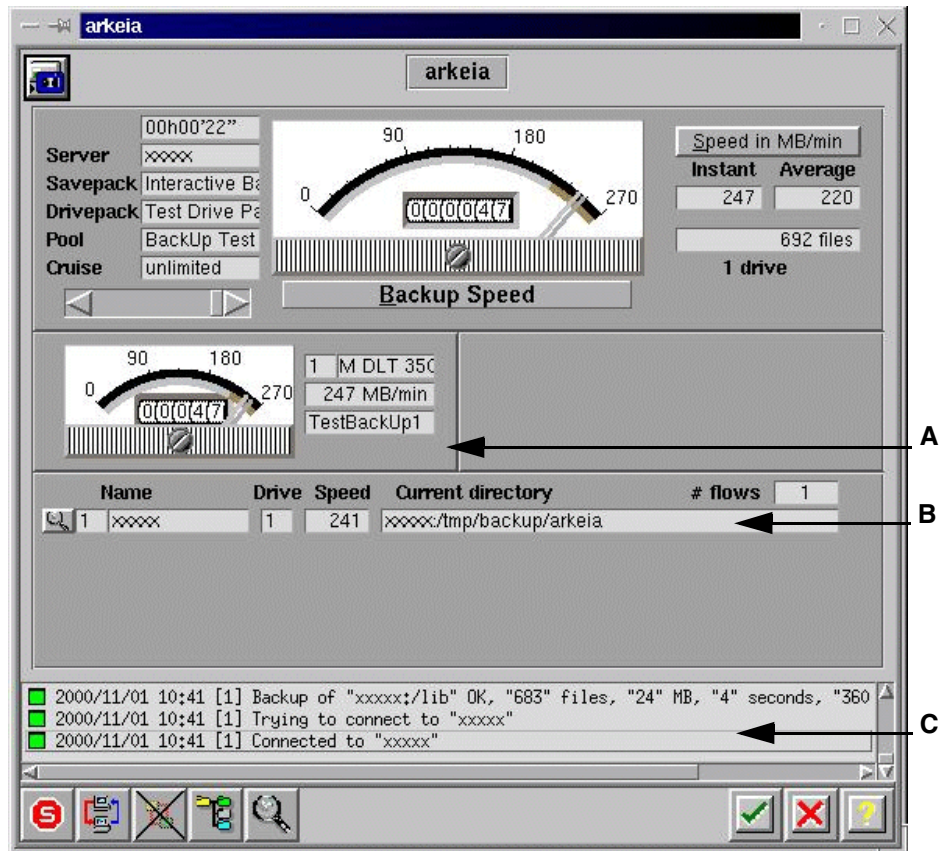


Figure 298. Main window during backup in progress

There are three areas in the window, marked **A**, **B** and **C** in Figure 298, which may require your attention:

In the area pointer **A** points to, you may sometimes see a push button labeled **OK**. Click this button when you have done the action, which was requested in the scroll list area **C**. In the line pointed to by **B**, you see the name of the file that actually is backed up.

You can leave this window by clicking the **OK** button. The backup process continues in the background.

If you want to connect again to this process or - as Arkeia calls it - job, go to Arkeia's main dialog window as shown in Figure 280. In this window you will see a box labeled either "No job running" or "List of jobs". If you see the text "List of jobs" and one or more lines under this box, move the cursor over the

line with the job you want to connect to and press the right mouse button. A pull-down menu as shown in Figure 299 appears.




Figure 299. Connect job pull-down menu

Move the cursor over the line with the action you will perform and click the left mouse button. The requested action will be performed.

If you chose **Stop job**, you are asked in a new dialog whether you really want to stop this job.

If you select **Connect job**, you will see a window similar to Figure 298 again.

16.3.4 Periodic Backup

To configure your scheme for periodic backups, press the periodic backup button or go to **Utilities>Periodic Backup** on the menu. You will see a window similar to Figure 300.




Figure 300. Periodic Backup window

To create a new entry for periodic backup, click the **new** button. You can now fill in the fields with the appropriate information. For more details, please consult the *Administrator's Manual*.

16.3.5 Restoration

To start restoration of data, click the restoration button or click **Restoration -> Restoration** on the menu. You will see a window like Figure 301.




Figure 301. Restoration start dialog

Clicking with the left mouse button over the checkbox beside an item toggles the status of item between selected/not selected. By double-clicking over a symbol for a complete system or a directory, you can navigate through the tree of information, that this backup contains. If you are ready with your selection, click the **OK** button and a window like Figure 302 appears, containing a list of the files or directories that will be restored.

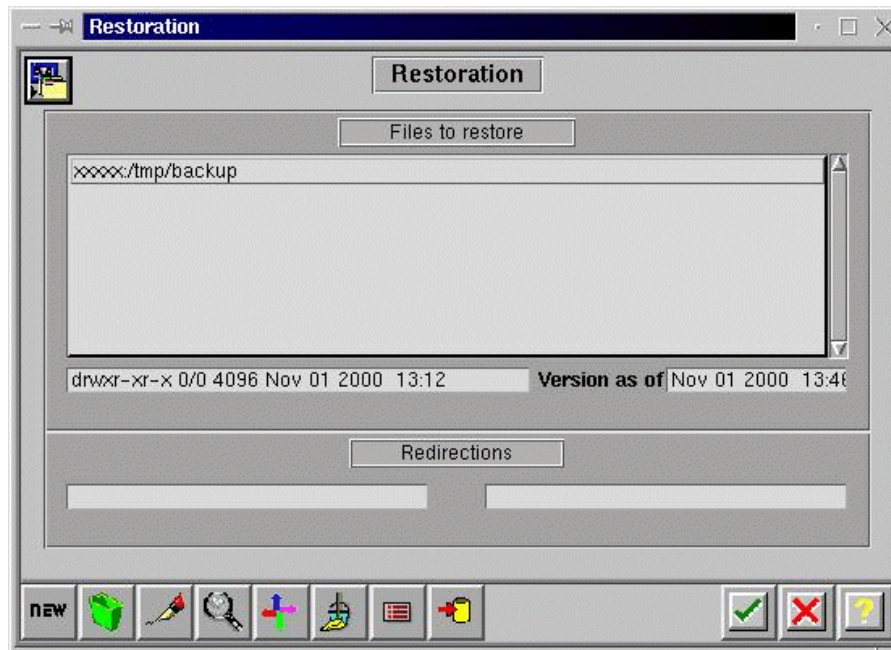


Figure 302. List of directories/files to store

Click the **OK** button in this window opens a new window, shown in Figure 303.



Figure 303. List of tapes used for restoration

You will see a list of the tape(s) that will be used during restoration. Click the **OK** button to proceed.

If the correct tape is already loaded to start the restoration with, you will see a window like Figure 304.

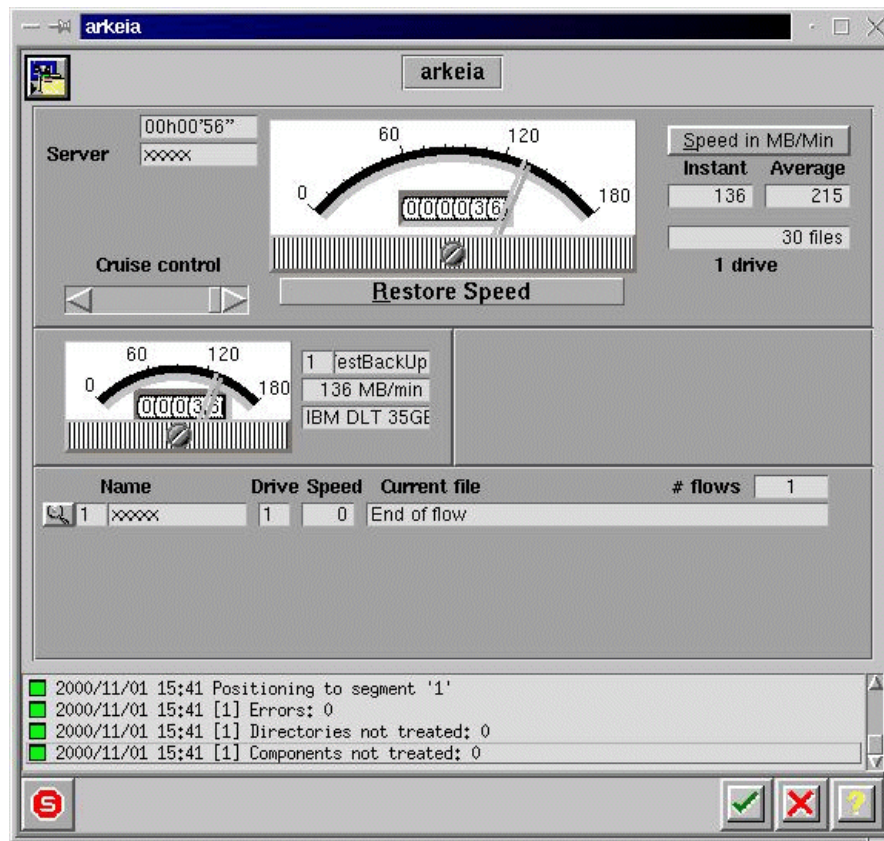


Figure 304. Restoration's main window

If the tape to start with must be mounted, a window like Figure 305 appears.

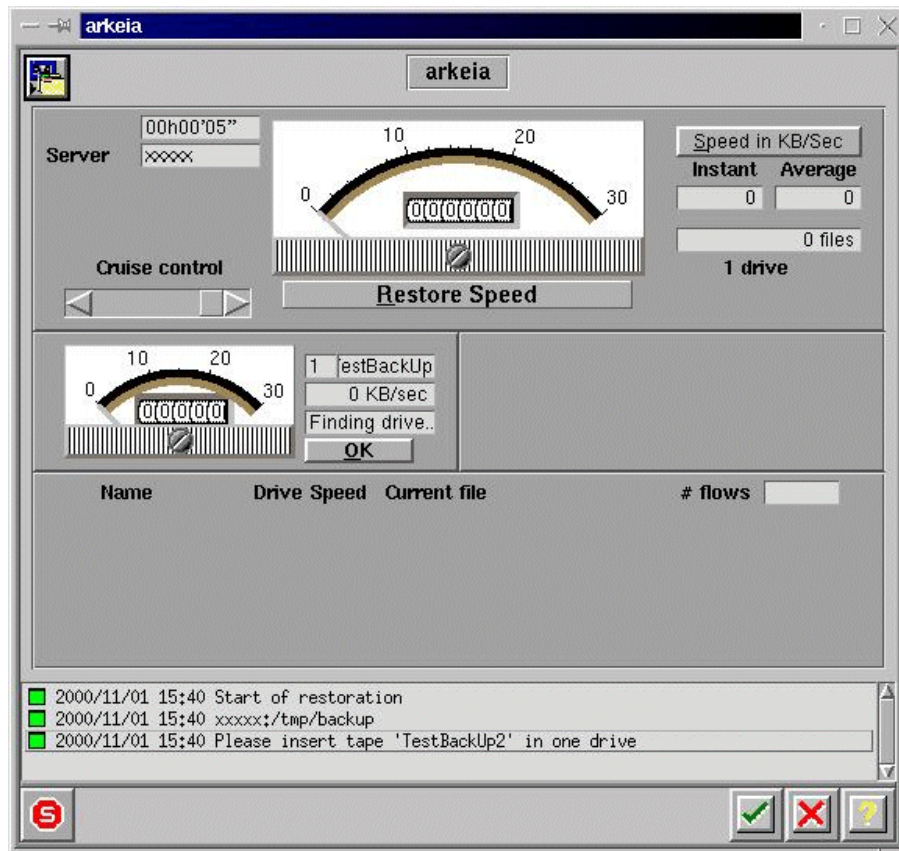


Figure 305. Window during restoration if manual intervention is required

Perform the action required and click **OK** to proceed. The appearance of the window changes. It is now like Figure 304.

16.3.6 Advanced features of Arkeia

For the advanced features of Arkeia, for example how to recycle or label tapes, please read the *Administrator's Manual*.

For more information, consult Arkeia's Web site at:

<http://www.arkeia.com>

Chapter 17. SNMP

SNMP stands for Simple Network Management Protocol. However, don't be fooled by the name SNMP can be powerful and complex. SNMP is an application-layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

The two primary components of an SNMP implementation are the SNMP agent and the Network Management Application.

- **SNMP agent**

An agent is a software component that resides on a managed device and collects management information. A managed device could be anything from a UPS, a router, or a computer.

- **Network Management Application**

The Network Management application can monitor and control devices on which an SNMP agent is running.

The three commands that are most commonly used in SNMP communication are read, write, and trap.

- **Read:** Used by the network management application to query SNMP agents for management information.
- **Write:** Used by the network management application to modify variables maintained by the SNMP agent.
- **Trap:** Used by SNMP agents to send alerts to Network management applications when defined thresholds are met or events occur.

17.1 Community Strings

The collection of management information that an agent is responsible for is called the Management Information Base or MIB. MIBs are organized hierarchically and are comprised of managed objects. These managed objects are identified by Object Identifiers or OIDs within the MIB hierarchy.

The MIB tree begins with the standards organizations: CCITT, ISO, and ISO-CCITT.

The objects that we will be looking at in our examples are located under the .iso.identified-organization.dod.internet.mgmt.mib-2 branch. This can also be referenced in short by the .1.3.6.1.2.1 descriptor.

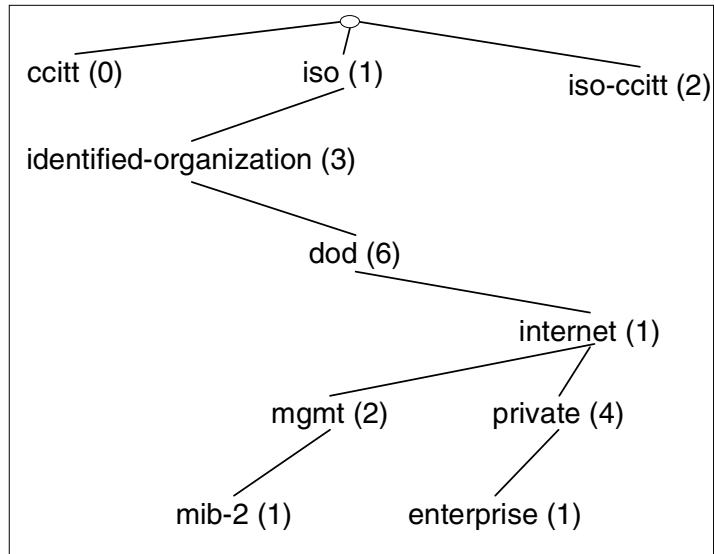


Figure 306. Partial MIB tree illustration

Under the mib-2 branch, we have several OIDs that contain basic system information. For example, from RFC1213-mib2.asn1:

- 1.3.6.1.2.1.1.1 - sysDescr

A textual description of the entity. This value should include the full name and version identification of the system's hardware type, software operating system, and networking software. It is mandatory that this only contain printable ASCII characters.
- 1.3.6.1.2.1.1.3 - sysUpTime

The time (in hundredths of a second) since the network management portion of the system was last re-initialized.
- 1.3.6.1.2.1.1.4 - sysContact

The textual identification of the contact person for this managed node, together with information on how to contact this person.
- 1.3.6.1.2.1.1.5 - sysName

An administratively assigned name for this managed node. By convention, this is the node's fully qualified domain name.
- 1.3.6.1.2.1.1.6 - sysLocation

The physical location of this node (e.g., 'telephone closet, 3rd floor').

Sure, it's nice to know the amount of time since a device last rebooted or the physical location, but let's look at something that may be more useful. Again under the MIB-2 branch, we have several OIDs that can be used to track network usage and also alert us to certain kinds of network errors. For example, for each interface:

- 1.3.6.1.2.1.2.2.1.8 - ifOperStatus

The current operational state of the interface (up, down, or testing). The testing(3) state indicates that no operational packets can be passed.

- 1.3.6.1.2.1.2.2.1.10 - ifInOctets

The total number of octets received on the interface, including framing characters.

- 1.3.6.1.2.1.2.2.1.16 - ifOutOctets

The total number of octets transmitted out of the interface, including framing characters.

- 1.3.6.1.2.1.2.2.1.20 - ifOutErrors

The number of outbound packets that could not be transmitted because of errors.

These examples contain general information about network interfaces. You have the ability to get much more specific the farther you drill down into the MIB sub-set.

Note

The .1.3.6.1.4.1 is where vendors such as IBM or Cisco locate customized objects for their products. There is practically no limit to the number of branches that are available. You should research vendor-specific MIB collections to get the most out of SNMP.

17.2 Why should I use SNMP?

SNMP is typically used to gauge network performance, find and solve network problems, and plan for network growth. However, you can also use SNMP to monitor vendor-specific hardware such as the current load on a UPS, the CPU utilization on routers, hubs, and computers, and even disk I/O and free space. The possibilities are endless.

You are not limited to predefined MIBs. Although it is beyond the scope of this document, you can compile your own MIBs. I urge you to take a look at the documentation included with the net-snmp packages.

Note

In an emergency, timing can be critical. Several commercial SNMP packages include paging software. If you wish to utilize paging and your software does not include it, take a look at the free application HylaFAX, included in SuSE.

17.3 Implementation on Linux

The SNMP package can be found in the n (networking) package selection section as:

```
snmp
```

The default configuration file installed with SNMP is fine for our purposes. To find more information about the configuration file syntax, look at the `snmpd.conf` man page.

To start `snmpd`, type:

```
rcsnmpd start
```

Congratulations. You have set up a basic SNMP agent.

To test our `snmp` implementation, we will use the `snmpget` command. This command queries SNMP agents on specified hosts for one or more OID values. The syntax is as follows:

```
snmpget HOST COMMUNITY OID
```

Try the following command:

```
snmpget localhost public .1.3.6.1.2.1.1.1.0
```

You should get a similar response to:

```
system.sysDescr.0 = OCTET STRING:"Linux Mail 2.2.16 i686"
```

The OID `.1.3.6.1.2.1.1.1` maps to the system description.

To see all of the available objects in our tree we will use the `snmpwalk` command. This command queries an entire tree instead of individual OIDs.

The basic syntax is the same as `snmpget` (although the two commands have several different options):

```
snmpwalk localhost public .1
```

With this command you “walk” the entire tree of OIDs that are available to you. You can use the `snmpwalk` and `snmpget` commands from a remote Linux host on the network and get the same result.

This is a very simplistic implementation of SNMP. Included with the SNMP package is a sample `snmpd.conf` file that includes methods for monitoring CPU utilization, disk space, and several other useful examples. With these packages also comes the ability to set traps to be sent to a specified host. The documentation included with `net-snmp` is quite thorough. Take a look.

17.3.1 MRTG

The Multi Router Traffic Grapher (MRTG) is a tool that utilizes SNMP to monitor the traffic load on network-links. MRTG generates HTML pages containing PNG images which provide a snapshot visual representation of this traffic. MRTG is an excellent example of what you can do with SNMP.




Figure 307. MRTG in action

MRTG requires Perl to be installed. If the Perl packages is not installed, simply mount the distribution media and install the following RPM:

```
perl
```

from the **a** (base system) section.

Although MRTG can be used without a Web server, we recommend you install the Apache Web server. See Chapter 7, “Apache and IBM HTTP Servers” on page 197.

Install the `mrtg` package from the **n**(networking) package selection section.

In our example we will create a configuration file to monitor and graph the network traffic on the localhost. We will use the `cfgmaker` tool to create the configuration file:

```
cd /usr/local/share/doc/packages/mrtg
./cfgmaker public@localhost > /etc/mrtg.conf
```

We need to edit the newly created `mrtg.conf`. In the section Global config options, uncomment and change the `WorkDir` entry for UNIX to the default HTML directory with the subdirectory `mrtg`. In our example we are using the Apache Web server, which serves Web pages from `/usr/local/httpd/htdocs`. Open the `mrtg.conf`, located in the following subdirectory:

```
WorkDir: /usr/local/httpd/htdocs/mrtg
```

Save the file.

Sample `mrtg.conf` file:

```
#####
# Description: Linux mail 2.2.16 i686
# Contact: jhaskins@ibm.com
# Location: Seattle, WA USA
#####

### Interface 2 >> Descr: 'eth0' | Name: '' | Ip: '192.168.1.1' | 05-44-99

Target[localhost_2]: 2:public@192.168.1.1
MaxBytes[localhost_2]: 96000
Title[localhost_2]: Traffic Analysis for 2 -- phu
PageTop[localhost_2]: <H1>Traffic Analysis for 2 -- phu</H1>
<TABLE>
  <TR><TD>System:</TD>      <TD>phu in Seattle</TD></TR>
  <TR><TD>Maintainer:</TD> <TD>jhaskins@uswest.net</TD></TR>
  <TR><TD>Description:</TD><TD>eth0 </TD></TR>
  <TR><TD>ifType:</TD>      <TD>ethernetCsmacd (6)</TD></TR>
  <TR><TD>ifName:</TD>      <TD></TD></TR>
  <TR><TD>Max Speed:</TD>   <TD>96.0 kBytes/s</TD></TR>
  <TR><TD>Ip:</TD>          <TD>10.0.0.254 ()</TD></TR>
</TABLE>
```

Now we can run MRTG against our config file:

```
[root@m10A bin]# mrtg /etc/mrtg.conf
Rateup WARNING: /usr/bin/rateup could not read the primary log file for
localhost.2
Rateup WARNING: /usr/bin/rateup The backup log file for localhost.2 was
invalid as well
```

```
Rateup WARNING: /usr/bin/rateup Can't remove localhost.2.old updating log
file
Rateup WARNING: /usr/bin/rateup Can't rename localhost.2.log to
localhost.2.old updating log file
```

The first two times you run MRTG against your config file you will get warnings. MRTG is looking for the old log files and can't find them because they don't exist yet.

Check to see if MRTG created the images and HTML pages.

```
ls /usr/local/httpd/htdocs/mrtg
```

```
localhost.2-day.png
localhost.2-month.png
localhost.2-week.png
localhost.2-year.png
localhost.2.html
localhost_2.log
localhost_2.old
```

Make sure your Web server is running and point your Web browser to:

```
http://localhost/mrtg/localhost.2.html
```

You should see a Web page similar to Figure 308 on page 358 (with less data of course):

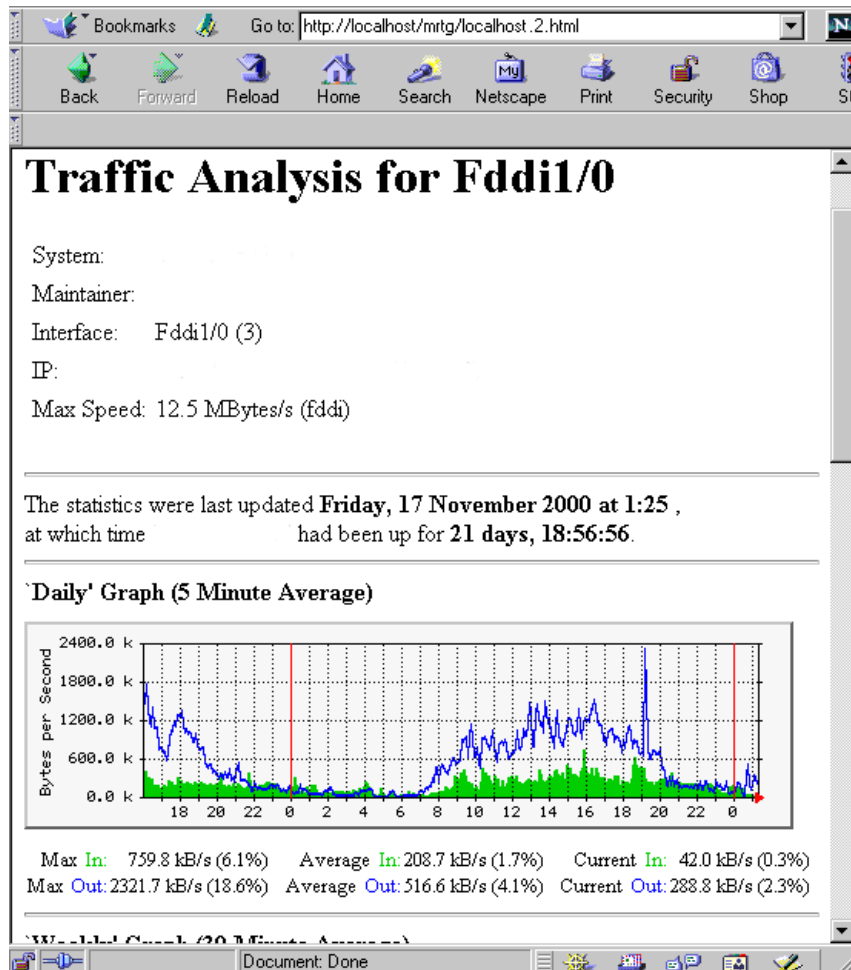


Figure 308. A sample MRTG page running for more than 24 hours on a FDDI interface

The final step is to automate the running of mrtg. Add the following entry to your crontab:

```
0/5 * * * * /usr/bin/mrtg /etc/mrtg.conf
```

Now MRTG will run every five minutes and update your Web site.

There is much more you can do with MRTG and SNMP. Check out the MRTG Web page for examples and documentation:

<http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>

17.3.2 Sources of additional information

Check out these Web sites for more ideas about how you can leverage the power of SNMP to help you manage your IT infrastructure.

Linux SNMP Network Management Tools:

<http://linas.org/linux/NMS.html>

OID assignments from the top node:

<http://www.alvestrand.no/harald/objectid/top.html>

Appendix A. RAID levels

This appendix has been included for the convenience of our readers who are unfamiliar with the disk subsystem technology known as RAID. We anticipate that this will be a small percentage of our readership, because RAID is an important technology that most people implementing business-critical IT systems probably know about. RAID is mentioned in many places throughout this book and a basic appreciation of its features and benefits will help you to understand why.

Even those who know about RAID already will be interested to hear about the new RAID-5E level supported by the latest IBM ServeRAID adapter.

A.1 What is RAID?

Although very commonly implemented using SCSI disks, RAID is independent of the specific disk technology being used. IBM Netfinity servers have RAID controllers that support SCSI, Fibre Channel, and SSA disk subsystems. In addition, Windows NT supports its own software-based RAID, although this is not often used, since much of the performance gained from having a dedicated hardware RAID controller is lost.

A typical RAID disk subsystem will have between two and six physical disks that are accessed by the processor by way of a specialized RAID controller adapter. The controller makes the array appear as a single large virtual disk to the processor. Because this disk has six completely independent head mechanisms for accessing data (in the case of a six-drive array), the potential for improved performance is immediately apparent. In the optimal situation, all six heads could be providing data to the system without the need for the time-consuming head-seeks to different areas of the disk that would be necessary were a single physical disk being used.

However, the primary intent of a RAID implementation is to prevent the system served by the array from being affected by critical hard disk failures. Several different implementations of RAID have been defined and are referred to as levels. Each level has different characteristics and these levels allow a choice to be made to best meet the cost, security, and performance desired. The three most common implementations are levels 0, 1, and 5. These are the levels available with all of IBM's disk subsystems supported by Netfinity servers, namely SCSI, SSA, and Fibre Channel. The Netfinity ServeRAID-3HB Ultra2 SCSI adapter introduces a new enhanced RAID-5 described in A.1.5, "RAID-5 enhanced" on page 369.

A.1.1 RAID-0

RAID-0, sometimes referred to as disk striping, is not really a RAID solution since there is no redundancy in the array at all. The disk controller merely stripes the data across the array so that a performance gain is achieved. This is illustrated in Figure 309:




Figure 309. RAID-0 implementation

It is common for a striped disk array to map data in blocks with a stripe size that is an integer multiple of real drive track capacity. For example, the IBM ServeRAID adapters allow stripe sizes of 8 KB, 16 KB, 32 KB or 64 KB, selectable during initialization of the array. Applications get better performance if their data I/O size matches the stripe size of the array, so it is recommended that you take this into consideration when defining your RAID sets.

Advantages:

- Performance improvement in many cases.
- All disk space available for data.

Disadvantages:

- No redundancy.

A.1.2 RAID-1 and RAID-1E

RAID-1, or disk mirroring, offers true redundancy. Each stripe is duplicated, or mirrored, on another disk in the array. In its simplest form, there are two disks where the second is a simple copy of the first. If the first disk fails then the second can be used without any loss of data. Some performance enhancement is achieved by reading data from both drives. Certain operating systems, including Windows NT, provide direct support for disk mirroring. There is a performance overhead, however, as the processor has to issue duplicate write commands. Hardware solutions where the controller handles the duplicate writes are preferred.

When more than two disks are available, the duplication scheme can be a little more complex to allow striping with disk mirroring, also known as Enhanced RAID-1. An example is shown in Figure 310:




Figure 310. RAID-1E implementation

As you can see, any one disk can be removed from the array without loss of information because each data stripe exists on two physical disks. The controller detects a failed disk and redirects requests for data from the failed drive to the drive containing the copy of the data. When a drive has failed, the replacement drive can be rebuilt using the data from the remaining drives in the array.

When a disk fails, there is only one copy of the data that was on the failed disk available to the system. The system has lost its redundancy, and if another disk fails, data loss is the result. To avoid this, failed disks should be replaced as soon as possible. The controller then rebuilds the data that was on the failed disk from the remaining drives and writes it to the new disk, restoring the redundancy.

To avoid having to manually replace a failed disk, the IBM Netfinity ServeRAID controllers implement *hot spare* disks. A hot spare disk is held idle until a failure occurs, at which point the controller immediately starts to rebuild the lost data onto the hot spare, minimizing the time when redundancy is lost. The controller continues to provide data to the system while the rebuild takes place.

When you replace the failed drive, its replacement becomes the array's new hot spare.

Advantages:

- Performance improvement in many cases.
- Redundancy. A drive can fail without loss of data.

Disadvantages:

- Cost. The logical disk has only half the capacity of the physical disks.

A.1.3 RAID-10

As we have seen, RAID-1 offers the potential for performance improvement as well as redundancy. RAID-10 is a variant of RAID-1 that effectively creates a mirror copy of a RAID-0 array.

In large disk subsystems that require, for example, two external storage enclosures, it would be beneficial to ensure that mirrored data exists in both units. This would allow an entire unit, including its power supply or connecting cables, to fail without interrupting operation. RAID-10 does just this by allowing one RAID-0 array to be contained in one of the enclosures and its mirror copy in the other. A diagram of a RAID-10 configuration is shown below:

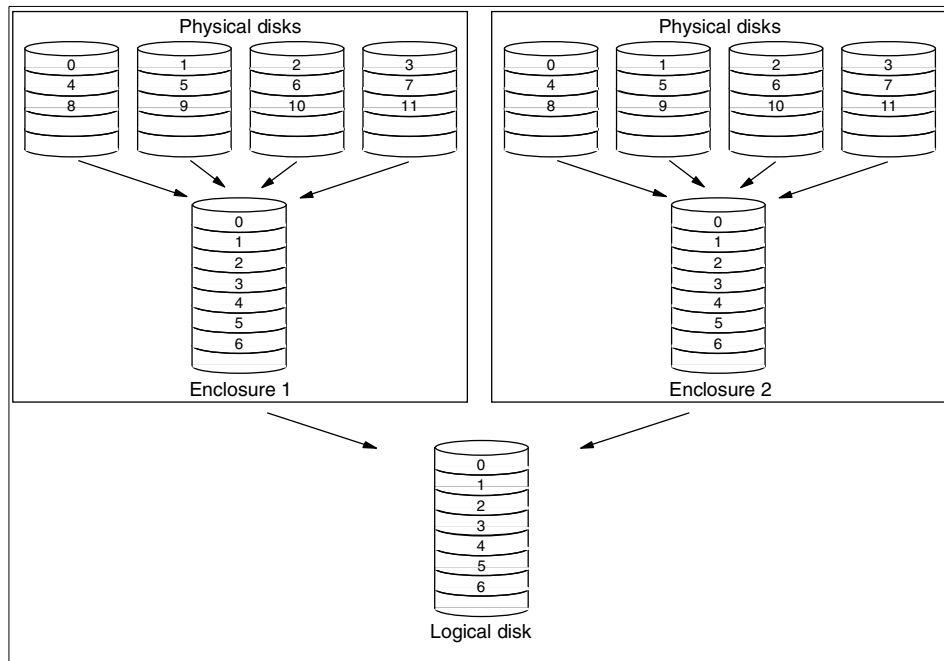


Figure 311. RAID-10 configuration

RAID-10 configurations are supported by the IBM Netfinity Fibre Channel RAID Controller Unit.

Advantages:

- Performance improvement in many cases.
- Redundancy. A drive can fail without loss of data.
- Provides fault tolerance for disk enclosures.

Disadvantages:

- Cost. The logical disk has only half the capacity of the physical disks.
- Slightly less flexible than RAID-1E (requires an even number of disks).

A.1.4 RAID-5

RAID-5 is one of the most capable and efficient ways of building redundancy into the disk subsystem. The way redundancy is implemented, capacity loss is equal to one of the drives in the array and data striping provides the read performance gains from RAID-0 and RAID-1. The principles behind RAID-5 are very simple and are closely related to the parity methods sometimes used for computer memory subsystems. In memory, the parity bit is formed by

evaluating the number of 1 bits in a single byte. For RAID-5, if we take the example of a four-drive array, three stripes of data are written to three of the drives and the bit-by-bit parity of the three stripes is written to the fourth drive.

As an example, we can look at the first byte of each stripe and see what this means for the parity stripe. Let us assume that the first byte of stripes 1, 2, and 3 are the letters A, B, and G respectively. The binary code for these characters is 01000001, 01000010 and 01000111 respectively.

We can now calculate the first byte of the parity block. Using the convention that an odd number of 1s in the data generates a 1 in the parity, the first parity byte is 01000100 (see Table 22). This is called Even Parity because there is always an even number of 1s if we look at the data and the parity together. Odd parity could have been chosen; the choice is of no importance as long as it is consistent.

Table 22. Generation of parity data for RAID-5

Disk 1 "A"	Disk 2 "B"	Disk 3 "G"	Disk 4 Parity
0	0	0	0
1	1	1	1
0	0	0	0
0	0	0	0
0	0	0	0
0	0	1	1
0	1	1	0
1	0	1	0

Calculating the parity for the second byte is performed using the same method, and so on. In this way, the entire parity stripe for the first three data stripes can be calculated and stored on the fourth disk.

The presence of parity information allows any disk to fail without loss of data.

In the above example, if drive 2 fails (with B as its first byte) there is enough information in the parity byte and the data on the remaining drives to reconstruct the missing data. The controller has to look at the data on the remaining drives and calculate what drive 2's data must have been to

maintain even parity. Because of this, a RAID-5 array with a failed drive can continue to provide the system with all the data from the failed drive.

Performance will suffer, of course, because the controller has to look at the data from all drives when a request is made to the failed one. However, that is better than losing the system completely. A RAID-5 array with a failed drive is said to be critical, since the loss of another drive will cause lost data. For this reason, the use of hot spare drives in a RAID-5 array is as important as in RAID-1.

The simplest implementation would always store the parity on disk 4 (in fact, this is the case in RAID-4, which is hardly ever implemented for the reason about to be explained). Disk reads are then serviced in much the same way as a level 0 array with three disks. However, writing to a RAID-5 array would then suffer from a performance bottleneck. Each write requires that both real data and parity data are updated. Therefore, the single parity disk would have to be written to every time any of the other disks were modified. To avoid this, the parity data is also striped, as shown in Figure 312, spreading the load across the entire array.

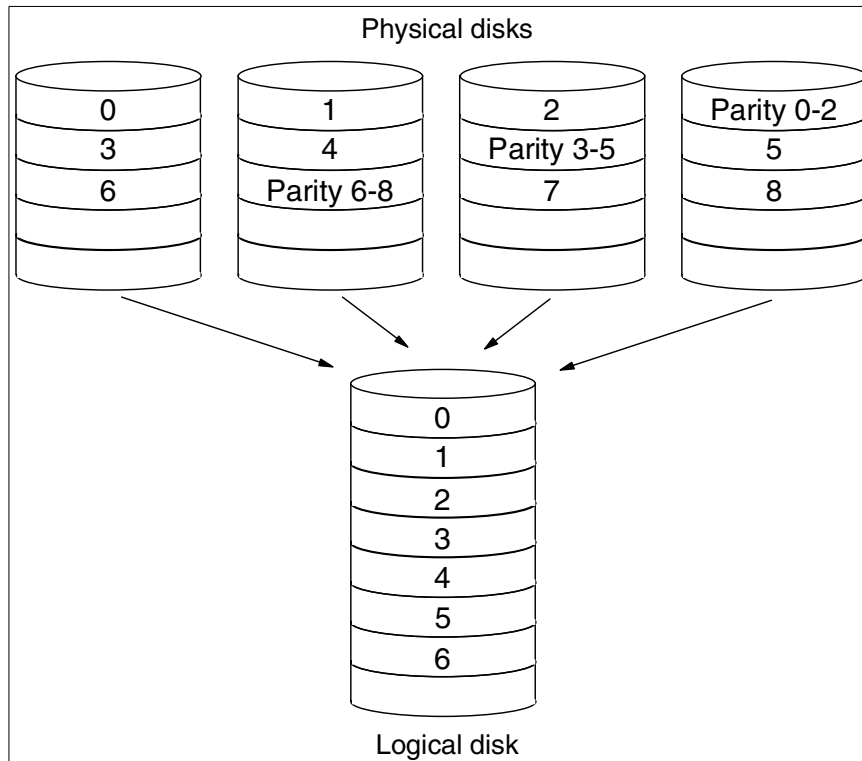


Figure 312. RAID-5 implementation

The consequence of having to update the parity information means that for every stripe written to the virtual disk, the controller has to read the old data from the stripe being updated and the associated parity stripe. Then the necessary changes to the parity stripe have to be calculated based on the old and the new data. All of this complexity is hidden from the processor, but the effect on the system is that writes are much slower than reads. This can be offset to a greater or lesser extent by the use of a cache on the RAID controller. The IBM ServeRAID controllers have cache as standard, which is used to hold the new data while the calculations are being performed. Meanwhile, the processor can continue as though the write has taken place. Battery backup options for the cache, available for some controllers, mean that data loss is kept to a minimum even if the controller fails with data still in the cache.

Advantages:

- Performance improvement in many cases.
- Redundancy. A drive can fail without loss of data.

- Storage overhead is equal to the size of only one drive.

Disadvantages:

- Overhead associated with writes can be detrimental to performance in applications where the write/read ratio is high. A controller cache can alleviate this.

A.1.5 RAID-5 enhanced

RAID-5 Enhanced (RAID-5E) puts hot spare drives to work to improve reliability and performance. A hot spare is normally inactive during array operation and is not used until a drive fails. By utilizing unallocated space on the drives in the array, a virtual distributed hot spare (DHS) can be created to improve reliability and performance. Figure 313 shows normal operation of a RAID-5E array. The data areas of the individual disks shown contain the application data and stripe parity data as for a normal RAID-5 array:

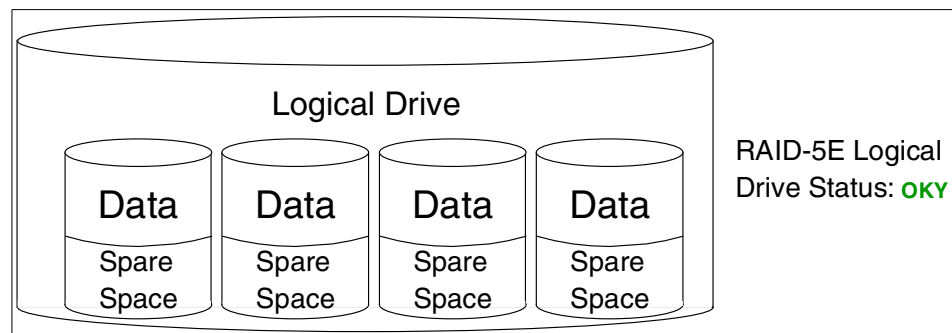


Figure 313. RAID-5E array: normal operation

In the event of a physical drive failing, its status will change to Defunct Disk Drive (DDD) and the ServeRAID adapter will start rearranging the data the disk contained into the spare space on the other drives in the array, provided there is enough space, of course.

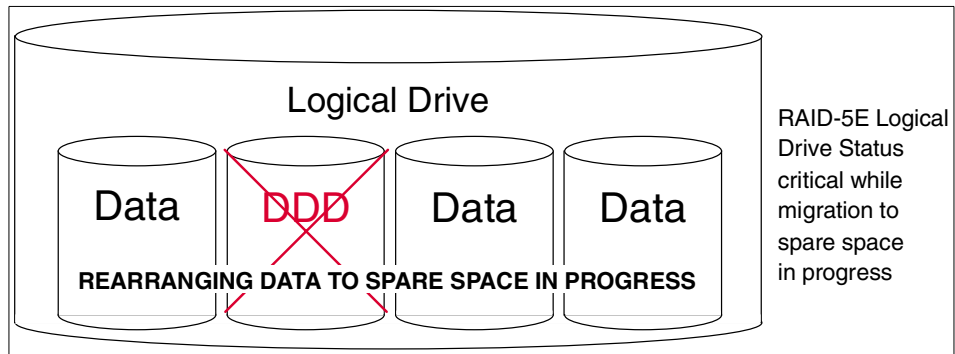


Figure 314. RAID-5E array: single physical disk failure

During the migration of data, the logical drive will be in a critical, nonredundant state. As soon as all the data is rearranged, the logical drive will be marked OKY (Okay) and have full redundancy again. This is illustrated in Figure 315.




Figure 315. RAID-5E array: data distributed throughout previous spare space

In the event of a second physical disk failure before the previously failed disk has been replaced, illustrated in Figure 316, normal RAID-5 procedures will be taken to provide service to the system through the checksum calculations described in A.1.4, “RAID-5” on page 365.




Figure 316. RAID-5E array: second physical disk failure

Advantages (compared to RAID-5):

- 15 - 20% performance improvement for smaller arrays with typical data transfer size.
- Protects data, even in the event of a two-drive failure.

Disadvantages:

- Migration time.

Design characteristics:

- One RAID-5E logical drive per array.
- Minimum of four physical drives in array configured for RAID-5E logical drive.

A.1.6 Orthogonal RAID-5

Orthogonal RAID-5 is an enhancement of RAID-5 in the sense that it is powered by more than one disk controller and hence improves both reliability and performance.

The performance of a disk subsystem depends on more than just the underlying performance of the disks. Multiple requests to one disk or across one adapter will typically take longer to satisfy than the same number of requests to multiple disks across multiple adapters.

In addition, the overall reliability of a standard RAID-5 system is dependent on the reliability of the one disk adapter to which all of the disks are connected. Orthogonal RAID-5 solves both of these concerns by grouping the disk arrays orthogonally to the disk adapters, SCSI buses, and power cables.

This would normally be implemented as a four-drive orthogonal RAID-5 array, where each disk would be connected to a different adapter and SCSI bus.

The result of this is that any one component of the disk subsystems, not just a disk drive, can fail with no loss of data and no interruption to system operation.

A.1.7 Performance

With different parameters affecting your RAID solution it is virtually impossible to find the perfect combination without measuring live throughput. Increasing redundancy also increases price and possibly lowers performance due to added overhead, which could be solved with more or faster controllers, again increasing the price.

As you can see in Figure 317 on page 373, speed is a significant issue when deciding on RAID level. The numbers shown in this figure and in Figure 318 on page 374 are based on benchmark testing performed by the IBM `@server` xSeries server development team. Specific systems may not show precisely the same performance ratios but the figures are representative of typical performance data.




Figure 317. Relative server performance versus RAID strategy

It is important to point out that the speed difference in Figure 317 is mainly due to the same number of drives being used for all tests. Generally, the more drives you use in your array, the faster it gets, but it also requires your RAID controller to be able to attach more drives when using RAID-1 or RAID-5 to get optimal performance.

Using the same number of drives:

- RAID-0 gives up to 50% more throughput than RAID-1.
- RAID-1 gives up to 50% more throughput than RAID-5.

The above test was done using a worst-case scenario with 50% reads and 50% writes. A high write/read ratio adversely affects the performance of RAID-1 and RAID-5 arrays, so throughput improves with a higher percentage of reads, which is generally more common in a real-world environment.

- While increasing the number of drives boosts performance, it also increases the price. Figure 318 on page 374 shows what happens with I/O throughput when we add drives to a RAID-0 array.




Figure 318. Adding drives to an array

Server throughput improves up to 50% when the number of drives is doubled for a RAID-0 and similar gains are shown for RAID-1 and RAID-5.

A.1.8 Recommendations

Before configuring your array you have to decide on a stripe size for the array. When configuring for maximum performance, Table 23 shows some rules of thumb:

Table 23. Recommended stripe configurations for ServeRAID adapters

Environment	Stripe size	Read-ahead
Groupware (Lotus Notes, Exchange)	16 KB	ON
Database Server (Oracle, SQL Server, DB/2)	16 KB	OFF
File Server (Windows NT 4.0, NetWare 4.1x)	16 KB	ON

Environment	Stripe size	Read-ahead
Web Server	8 KB	OFF
Other	8 KB	ON

A.1.9 Summary

RAID is an excellent and proven technology for protecting your data against the possibility of hard disk failure. IBM has a range of RAID controllers that bring the benefits of the technology to @server xSeries and Netfinity servers. As Intel-based servers become more and more critical to customers' businesses, they are demanding the reliability provided by RAID.

Here is a quick summary of the different RAID levels we have covered in this appendix:

RAID-0: Block interleave data striping without parity

- Best performance of all RAID levels
- Drive seek times and latencies effectively reduced by parallel operation
- Significantly outperforms single large disk

RAID-1: Disk mirroring

- Fast and reliable but requires 100% disk space overhead
- Two copies of data maintained
- No performance degradation with a single disk failure
- Writes are slower than a single disk, reads are quicker

RAID-1E: Data stripe mirroring

- All the benefits of RAID-1
- Provides mirroring with an odd number of drives

RAID-10: Mirrored RAID-0 arrays

- All the benefits of RAID-1
- Can provide fault tolerance for entire storage enclosures

RAID-5: Block interleave data striping with distributed parity

- Best for random transactions
- Poor for large sequential reads if request is larger than block size
- Block size is the key to performance; must be larger than typical request size

- Performance degrades in recovery mode, that is, when a single drive has failed

RAID-5E: RAID-5 with distributed hot spare

- All the benefits of RAID-5
- 15 - 20% performance improvement for smaller arrays
- Protects data, even in the event of a two-drive failure

Orthogonal RAID-5: RAID-5 with multiple orthogonal disk adapters

- All the benefits of RAID-5
- Improved performance (due to load being spread across disk adapters)
- Improved reliability due to redundancy of disk adapters and disks

Table 24 gives you a summary of RAID performance characteristics:

Table 24. Summary of RAID performance characteristics

RAID level	Capacity	Large transfers	I/O rate	Data availability
RAID-0	Excellent	Very Good	Very Good	Poor ¹
RAID-1/1E	Moderate	Good	Good	Good
RAID-10	Moderate	Good	Good	Very Good
RAID-5	Very Good	Very Good	Good	Good
RAID-5E	Very Good	Very Good	Good to Very Good	Very Good
Orthogonal RAID-5	Very Good	Very Good	Good	Very Good

¹ Availability = MTBF of one disk divided by the number of disks in the array

If you want to learn more about RAID, the RAID Advisory Board, of which IBM is an active member, exists to standardize terminology and provide information about RAID technology. Its Web site can be found at the following URL:

<http://www.raid-advisory.com/>

Appendix B. Sample smb.conf SAMBA configuration file

```
# This is the main Samba configuration file. You should read the
# smb.conf(5) manual page in order to understand the options listed
# here. Samba has a huge number of configurable options (perhaps
# too many!) most of which are not shown in this example
#
# Any line which starts with a ; (semi-colon) or a # (hash)
# is a comment and is ignored. In this example we will use a #
# for commentry and a ; for parts of the config file that you
# may wish to enable
#
# NOTE: Whenever you modify this file you should run the command "testparm"
# to check that you have not many any basic syntactic errors.
#
#===== Global Settings =====
[global]

# workgroup = NT-Domain-Name or Workgroup-Name
# workgroup = LINUXRLZ

# server string is the equivalent of the NT Description field
# server string = Samba Server on Caldera OpenLinux

# This option is important for security. It allows you to restrict
# connections to machines which are on your local network. The
# following example restricts access to two C class networks and
# the "loopback" interface. For more examples of the syntax see
# the smb.conf man page
; hosts allow = 192.168.1. 192.168.2. 127.

# If you want to automatically load your printer list rather
# than setting them up individually then you'll need this
# load printers = yes

# you may wish to override the location of the printcap file
; printcap name = /etc/printcap

# It should not be necessary to specify the print system type unless
# it is non-standard. Currently supported print systems include:
# bsd, sysv, plp, lprmng, aix, hpux, qnx
# printing = lprmng

# Uncomment this if you want a guest account, you must add this to /etc/passwd
```

```

# otherwise the user "nobody" is used
; guest account = pcguest

# this tells Samba to use a separate log file for each machine
# that connects
; log file = /var/log/samba.d/smb.%m

# Put a capping on the size of the log files (in Kb).
    max log size = 50

# Security mode. Most people will want user level security. See
# security_level.txt for details.
    security = user
# Use password server option only with security = server
; password server = <NT-Server-Name>

# Password Level allows matching of _n_ characters of the password for
# all combinations of upper and lower case.
; password level = 8
; username level = 8

# You may wish to use password encryption. Please read
# ENCRYPTION.txt, Win95.txt and WinNT.txt in the Samba documentation.
# Do not enable this option unless you have read those documents
    encrypt passwords = yes
    smb passwd file = /etc/samba.d/smbpasswd

# The following are needed to allow password changing from Windows to
# update the Linux sytsem password also.
# NOTE: Use these with 'encrypt passwords' and 'smb passwd file' above.
# NOTE2: You do NOT need these to allow workstations to change only
#         the encrypted SMB passwords. They allow the Unix password
#         to be kept in sync with the SMB password.
; unix password sync = Yes
; passwd program = /usr/bin/passwd %u
; passwd chat = *New*UNIX*password* %n\n *ReType*new*UNIX*password* %n\n
*passwd:*all*authentication*tokens*updated*successfully*

# Unix users can map to different SMB User names
; username map = /etc/samba.d/smbusers

# Using the following line enables you to customise your configuration
# on a per machine basis. The %m gets replaced with the netbios name
# of the machine that is connecting
; include = /etc/samba.d/smb.conf.%m

```



```

# Most people will find that this option gives better performance.
# See speed.txt and the manual pages for details
    socket options = TCP_NODELAY

# Configure Samba to use multiple interfaces
# If you have multiple network interfaces then you must list them
# here. See the man page for details.
;   interfaces = 192.168.12.2/24 192.168.13.2/24

# Configure remote browse list synchronisation here
# request announcement to, or browse list sync from:
# a specific host or from / to a whole subnet (see below)
;   remote browse sync = 192.168.3.25 192.168.5.255
# Cause this host to announce itself to local subnets here
;   remote announce = 192.168.1.255 192.168.2.44

# Browser Control Options:
# set local master to no if you don't want Samba to become a master
# browser on your network. Otherwise the normal election rules apply
;   local master = no

# OS Level determines the precedence of this server in master browser
# elections. The default value should be reasonable
;   os level = 33

# Domain Master specifies Samba to be the Domain Master Browser. This
# allows Samba to collate browse lists between subnets. Don't use this
# if you already have a Windows NT domain controller doing this job
;   domain master = yes

# Preferred Master causes Samba to force a local browser election on startup
# and gives it a slightly higher chance of winning the election
;   preferred master = yes

# Use only if you have an NT server on your network that has been
# configured at install time to be a primary domain controller.
;   domain controller = <NT-Domain-Controller-SMBName>

# Enable this if you want Samba to be a domain logon server for
# Windows95 workstations.
;   domain logons = yes

# if you enable domain logons then you may want a per-machine or
# per user logon script

```

```

# run a specific logon batch file per workstation (machine)
; logon script = %m.bat
# run a specific logon batch file per username
; logon script = %U.bat

# Where to store roving profiles (only for Win95 and WinNT)
# %L substitutes for this servers netbios name, %U is username
# You must uncomment the [Profiles] share below
; logon path = \\%L\Profiles\%U

# All NetBIOS names must be resolved to IP Addresses
# 'Name Resolve Order' allows the named resolution mechanism to be specified
# the default order is "host lmhosts wins bcast". "host" means use the unix
# system gethostbyname() function call that will use either /etc/hosts OR
# DNS or NIS depending on the settings of /etc/host.config, /etc/nsswitch.conf
# and the /etc/resolv.conf file. "host" therefore is system configuration
# dependant. This parameter is most often of use to prevent DNS lookups
# in order to resolve NetBIOS names to IP Addresses. Use with care!
# The example below excludes use of name resolution for machines that are NOT
# on the local network segment
# - OR - are not deliberately to be known via lmhosts or via WINS.
; name resolve order = wins lmhosts bcast

# Windows Internet Name Serving Support Section:
# WINS Support - Tells the NMBD component of Samba to enable it's WINS Server
; wins support = yes

# WINS Server - Tells the NMBD components of Samba to be a WINS Client
# Note: Samba can be either a WINS Server, or a WINS Client, but NOT both
; wins server = w.x.y.z

# WINS Proxy - Tells Samba to answer name resolution queries on
# behalf of a non WINS capable client, for this to work there must be
# at least one WINS Server on the network. The default is NO.
; wins proxy = yes

# DNS Proxy - tells Samba whether or not to try to resolve NetBIOS names
# via DNS nslookups. The built-in default for versions 1.9.17 is yes,
# this has been changed in version 1.9.18 to no.
dns proxy = no

# Case Preservation can be handy - system default is _no_
# NOTE: These can be set on a per share basis
; preserve case = no
; short preserve case = no

```

```

# Default case is normally upper case for all DOS files
; default case = lower
# Be very careful with case sensitivity - it can break things!
; case sensitive = no

#===== Share Definitions =====
[homes]
    comment = Home Directories
; this gives access to a 'Public' sub-directory in each user's home...
; (it is named 'public' as it is intended to be used by other sharing
; technologies (like NetWare, appletalk) too and may get disclosed due
; to weak protocols! -- hmm, are there less secure protocols than NFS? :)
    path = %H
    valid users = %S
%    only user = yes
    browseable = no
    writable = yes
    create mask = 0750

# Un-comment the following and create the netlogon directory for Domain Logons
; [netlogon]
;    comment = Samba Network Logon Service
;    path = /home/samba/netlogon
;    guest ok = yes
;    writable = no
;    share modes = no

# Un-comment the following to provide a specific roving profile share
# the default is to use the user's home directory
; [Profiles]
;    path = /home/samba/profiles
;    browseable = no
;    guest ok = yes

# NOTE: If you have a BSD-style print system there is no need to
# specifically define each individual printer
[printers]
    comment = All Printers
    path = /var/spool/samba
    browseable = no
# Set public = yes to allow user 'guest account' to print
    guest ok = no
    writable = no

```

```

printable = yes
create mask = 0700

# A publicly accessible directory, but read only, except for people in
# the "users" group
[public]
    comment = Public Stuff
    path = /home/public
    browseable = yes
    public = yes
    writable = yes
    printable = no
# access may be controlled by these options
; read list = user1, user2, @group
; valid users = user1, user3
    write list = @users

# Other examples.
#
# This one is useful for people to share files, BUT
# access to '/tmp' or '/var/tmp' should *not* be given lightly,
# as this can (still) pose a security threat!
# Better use a dedicate sub-directory to /(var/)tmp or something
# like a [public] share!
;[tmp]
;    comment = Temporary file space
;    path = /tmp
;    read only = no
;    public = yes

# A private printer, usable only by fred. Spool data will be placed in fred's
# home directory. Note that fred must have write access to the spool directory,
# wherever it is.
;[fredsprn]
;    comment = Fred's Printer
;    valid users = fred
;    path = /homes/fred
;    printer = freds_printer
;    public = no
;    writable = no
;    printable = yes

# A private directory, usable only by fred. Note that fred requires write
# access to the directory.
;[fredsdir]

```

```

; comment = Fred's Service
; path = /usr/somewhere/private
; valid users = fred
; public = no
; writable = yes
; printable = no

# a service which has a different directory for each machine that connects
# this allows you to tailor configurations to incoming machines. You could
# also use the %u option to tailor it by user name.
# The %m gets replaced with the machine name that is connecting.
;[pchome]
; comment = PC Directories
; path = /usr/pc/%m
; public = no
; writable = yes

# A publicly accessible directory, read/write to all users. Note that all files
# created in the directory by users will be owned by the default user, so
# any user with access can delete any other user's files. Obviously this
# directory must be writable by the default user. Another user could of course
# be specified, in which case all files would be owned by that user instead.
;[public]
; path = /usr/somewhere/else/public
; public = yes
; only guest = yes
; writable = yes
; printable = no

# The following two entries demonstrate how to share a directory so that two
# users can place files there that will be owned by the specific users. In this
# setup, the directory should be writable by both users and should have the
# sticky bit set on it to prevent abuse. Obviously this could be extended to
# as many users as required.
;[myshare]
; comment = Mary's and Fred's stuff
; path = /usr/somewhere/shared
; valid users = mary fred
; public = no
; writable = yes
; printable = no
; create mask = 0765

```

Appendix C. Special notices

This publication is intended to help customers, business partners and IBM employees implement SuSE Linux. The information in this publication is not intended as the specification of any programming interfaces that are provided by SuSE Linux. See the PUBLICATIONS section of the IBM Programming for information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers

attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

IBM ®	Redbooks
DB2	Redbooks Logo 
Home Director	WebSphere
@server	XT
Netfinity	400
NetVista	Notes
PS/2	xSeries
ServeRAID	SP

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere., The Power To Manage., Anything. Anywhere., TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Københavns Sommer - Tivoli A/S.

SuSE and its logo are registered trademarks of SuSE AG. Linux is a trademark of Linus Torvalds.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 IBM Redbooks

For information on ordering these publications see “How to get IBM Redbooks” on page 393.

- *Linux for WebSphere and DB2 Servers*, SG24-5850
- *Red Hat Linux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5853
- *Caldera OpenLinux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5861
- *TurboLinux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5862

D.2 IBM Redbooks collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at ibm.com/redbooks for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
IBM System/390 Redbooks Collection	SK2T-2177
IBM Networking Redbooks Collection	SK2T-6022
IBM Transaction Processing and Data Management Redbooks Collection	SK2T-8038
IBM Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
IBM AS/400 Redbooks Collection	SK2T-2849
IBM Netfinity Hardware and Software Redbooks Collection	SK2T-8046
IBM RS/6000 Redbooks Collection	SK2T-8043
IBM Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

D.3 Other resources

These publications are also relevant as further information sources:

- *Understanding and Deploying LDAP Directory Services*, by Timothy Howes, Mark Smith, and Gordon Good, ISBN: 1578700701

- *Using Samba by Robert Eckstein, David Collier-Brown and Peter Kelly*, published by O'Reilly, available online at:
<http://www.oreilly.com/catalog/samba/chapter/book/index.html>
- *The Linux NIS (YP)/NYS/NIS+ HOWTO* by Thorsten Kakuk, found at
<http://metalab.unc.edu/pub/Linux/docs/HOWTO/NIS-HOWTO>.
- *Managing NFS and NIS*, by Hal Stern, ISBN 0937175757
- "Don't make me LDAP you - Lightweight Directory Access Protocol: What it is, why you want it", available from the LinuxWorld Web site at:
<http://www.linuxdoc.org/HOWTO/LDAP-HOWTO.html>.

D.4 Referenced Web sites

These Web sites are also relevant as further information sources:

- <http://www.suse.com>
- <http://www.redbooks.ibm.com>
- <http://www.fsf.org>
- <http://www.ibm.com/linux/>
- <http://www.redhat.com>.
- <http://www.developer.ibm.com/welcome/netfinity/serveraid.html>
- <http://www.linuxtr.net>
- <http://www.pc.ibm.com/support>
- <http://www.rpm.org>
- <http://www.solucorp.qc.ca/linuxconf>
- <http://www.linuxdoc.org>
- <http://www.linuxdoc.org/HOWTO/DNS-HOWTO.html>
- <http://www.samba.org>
- <http://www.netcraft.com/survey/>
- <http://www.apache.org>
- <http://www-4.ibm.com/software/webrowsers/httpservers/>
- http://www-4.ibm.com/software/webrowsers/httpservers/doc/v136/readme_httpserver.htm
- <http://www.apache.org/docs/misc/perf-tuning.html>
- <http://www-4.ibm.com/software/webrowsers/httpservers/download.html>
- <http://www.rustcorp.com/linux/ipchains>
- <http://www.sendmail.org>
- <http://www.dcs.qmw.ac.uk/~williams/>
- <http://www.metalab.unc.edu/pub/Linux/docs/HOWTO/NIS-HOWTO>

- <http://www.openldap.org/>
- <http://tune.linux.com>
- <http://www.tunelinux.com>
- <http://www.linux-mandrake.com/lothar/>
- <http://www.textuality.com/bonnie/>
- <http://www.netperf.org/netperf/NetperfPage.html>
- <http://www.estinc.com/>
- <http://www.microlite.com>
- <http://www.raid-advisory.com/>
- <http://www.raid-advisory.com>
- <http://www.elink.ibm.link.ibm.com/pbl/pbl>
- <http://w3.itso.ibm.com>
- <http://w3.ibm.com>

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** ibm.com/redbooks

Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at <http://w3.ibm.com/> for redbook, residency, and workshop announcements.

Index

A

AMD chipset 13
Apache 197
 features 197
 installation 199
 performance tips 208
ash shell 120

B

backup 279
backup and restore utility 279
bash shell 120
Beowulf 275
 packages 276
bind4 185
bind8 185
BM HTTP Server
 Administration Server 204
BRU 283
 basic backup 281
 basic restore 281
 commands 281
 installation 279
 restore 287
 scheduling 286

C

community strings 351
csh shell 121

D

DAP 248
DARPA 231
DHCP 38
disk striping 362
disk subsystem
 See also RAID
 RAID performance 372
DNS 183, 184, 185
 configuration 185
 installation 185
 YaST 186
dynamic host configuration protocol (DHCP) 221
 implementation 221

dynamic name server (DNS) 221

E

Ethernet 128

F

file/print 155
FQDN 183
FTP 15

H

hdparm 272

I

IBM HTTP Server 197, 198
 features 198
 installation 201
 performance tips 209
IBM Netfinity 3000 5
IBM Netfinity 3500 M10 5
IBM Netfinity servers 6
IBM xSeries 2
IETF 248
installing Linux
 ServeRAID 69
Intel 13
ISO 248

K

KDE 14, 265
ksh shell 120
KTop 265

L

LDAP 198
Lightweight Directory Access Protocol (LDAP) 247,
249
 ldap.conf file 250
 LDIF 249
 nsswitch.conf file 251
 PAM 251
 slapd.conf file 250
LILO 33, 34, 36
Linux commands 149

Linuxrc 8, 11, 13
locate 149

M

main.cf 224
master.cf 224
Microlite
 BackupEDGE 288, 290
 backup 292
 features 288, 324
 incremental backup 298
 installation 289
 master backup 298
 master restore 301
 restore 296
 schedule backup 302
 tape device 305
 RecoverEDGE 314
 boot disks 315
 features 314
 total crash 322
Midnight Commander (MC) 43
Multi Router Traffic Grapher (MRTG) 355

N

name resolution 184
netperf 273
network configuration 127
Network File System (NFS) 231, 232
 access data 233
 remote access 235
Network Information System (NIS) 237
 installation 238
 NIS+ 238
 server 239
network management application 351
NFS 15, 39
NIS 39
 installation 237

P

packet filtering with IP chains 211
 additional information 220
 dial-up Internet connection 211
 FTP masquerading 215
 gateway 211, 216
 checksum 216

 demasquerade 216
 forward chain 217
 input chain 216
 lo interface 216
 local process 216
 output chain 217
 routing decision 216
 sanity 216
IP chains 217
NAT 211
network configuration 212
requirements 212
performance
 of RAID subsystems 372
performance tools
 Powertweak 258
performance tools in Linux 257
Perl 355
Personal Systems Reference sheets 6
PHP 197
Postfix 223
Postfix configuration file 225
Postfix errors 229
POVray 275
pstree 265
PVMPOVray 275

R

RAID
 described 361
 level 0 (RAID-0) 362
 level 1 (RAID-1, RAID-1E) 363
 level 10 (RAID-10) 364
 level 5 (RAID-5) 365
 level 5 enhanced (RAID-5E) 369
 orthogonal RAID-5 371
 performance 372
 RAID Advisory Board 376
 recommendations 374
 software-based 361
 summary of RAID levels 375
 support for two disk failures 369
 supported disk technologies 361
recovery 279
Remote Procedure Call (RPC) 231
RFC 248
RPM 111, 117

S

- Samba 155
 - configuration 156
 - global settings 156
 - installation 155
 - NetBIOS 156
 - printer shares 163
 - shares 160
 - start 164
 - stop 164
 - SWAT 164
 - logon 166
 - WINS 155
- SaX 103, 104, 110
 - keyboard configuration 107
 - monitor configuration 109
 - mouse configuration 104, 105
 - screen selection 110
 - server selection 108
- Secure SHell (SSH) 193
 - configuration 193
- Sendmail 223
- ServeRAID 69
 - configuration 69
 - driver 69
 - hot swap rebuild 81
 - ipssend 72
 - commands 73
 - devinfo 79
 - getconfig 73
 - getstatus 78
 - hsrebuild 81
 - rebuild 85
 - setstate 82
 - synch 84
 - unattended 84
 - rebuild drive 85
 - replace drive 86
 - RPM 72
 - synchronize logical drives 84
 - unattended mode 84
 - utility 70
- sh shell 120
- simple network management protocol (SNMP) 351, 353, 359
- SNMP 198
- SNMP agent 351
- SSL 198
- SuSE
 - bug fixes 5
 - FTP server 5
 - SuSE Linux 6.3 185
- SuSE installation 5, 9
 - Adaptec 12
 - end / reboot 11
 - hardware considerations 5, 6
 - CD-ROM 5
 - CPU 6
 - display adapter 6
 - hard drives 5
 - monitor 6
 - mouse 6
 - network card 6
 - RAM 6
 - SCSI adapter 5
 - kernel modules 11
 - keyboard 11
 - language 10
 - LILO 34
 - mouse 42
 - network 13
 - network configuration 37
 - package installation 29, 32
 - partitioning 16, 17, 19, 21
 - root password 41
 - Sendmail 40
 - settings 11
 - start installation / system 11
 - swap 22, 24
 - system information 11
 - time zone 36
 - YaST 9
- SWAT 164
 - global settings 166, 168
 - logon 166
 - printers 175
 - restart Samba 175
 - Samba passwords 180
 - Samba status 179
 - shares 169
- system administration 111
 - command line tools 122, 123, 125
 - group administration 118
 - RPM 111
 - software packages 111
 - series 113
 - using RPM 117
 - user administration 118

T

tcsh shell 120
token-ring 128
top 264

U

updatedb 150

V

VESA 109
vmstat 272

W

Web site
RAID Advisory Board 376

X

X.500 247, 248
XFree86 14
xSeries brand 2
X-Windows 42

Y

YaST 8, 9, 15, 18, 21, 24, 41, 111, 118, 120
group administration 118
network configuration 127, 131
system administration 126
user administration 118

Z

zsh shell 121

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at ibm.com/redbooks
- Fax this form to: USA International Access Code + 1 845 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number	SG24-5863-01
Redbook Title	SuSE Linux Integration Guide for IBM @server for xSeries and Netfinity
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	<input type="radio"/> Very Good <input type="radio"/> Good <input type="radio"/> Average <input type="radio"/> Poor
Please identify yourself as belonging to one of the following groups:	<input type="radio"/> Customer <input type="radio"/> Business Partner <input type="radio"/> Solution Developer <input type="radio"/> IBM, Lotus or Tivoli Employee <input type="radio"/> None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	<input type="checkbox"/> Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. ibm.com/privacy/yourprivacy/



Redbooks

SUSE Linux Integration Guide for IBM @server for xSeries and Netfinity

(0.5" spine)

0.475" <-> 0.875"

250 <-> 459 pages



Redbooks

SuSE Linux Integration Guide for IBM **@server** for xSeries and Netfinity

The complete guide to running SuSE Linux on xSeries and Netfinity

Netfinity server-specific coverage you can't find anywhere else, including ServeRAID configuration

Plan, configure, and install key services, step-by-step: Samba, Apache, Postfix, DNS, DHCP, LDAP and more

Here's all the information you need to maximize SuSE Linux performance and reliability on the IBM state-of-the-art **@server** xSeries and Netfinity server platforms. In this book, a team of IBM's top Linux experts presents start-to-finish, Netfinity server-specific coverage of SuSE Linux 7.0 Professional deployment and system administration throughout the entire system life cycle!

You will get running fast with expert step-by-step preparation and installation techniques: updating your BIOS and firmware, making the CD-ROM bootable, preparing SCSI devices, partitioning, configuration, X-Windows setup, deploying IBM ServeRAID in SuSE Linux environments, and much more.

Next, you'll master all the key techniques of day-to-day SuSE Linux system administration, including backup and recovery, Internet and e-mail connectivity, DNS/DHCP name services, and using SuSE Linux with Samba as a world-class file/print server for Windows workstations. IBM-tested, proven, and crystal clear, this is the one essential book for everyone running SuSE Linux on xSeries and Netfinity.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-5863-01

ISBN 0739419869