Caldera OpenLinux Integration Guide

for IBM eserver xSeries and Netfinity

- The complete guide to running Caldera OpenLinux 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, sendmail, DNS, DHCP, LDAP, and more

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

Redbooks

ibm.com/redbooks

SG24-5861-01

International Technical Support Organization

Caldera OpenLinux Integration Guide for IBM @server xSeries and Netfinity

February 2001

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix H, “Special notices” on page 403.

Second Edition (February 2001)

This edition applies to preparing to or installing Caldera OpenLinux eServer 2.3 on IBM @server xSeries and Netfinity.

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.

Contents

Preface ix
The team that wrote this redbook x
Comments welcome xii

Chapter 1. Introduction 1
1.1 The IBM commitment to Linux 1
1.2 Caldera Systems, Inc. 2
1.3 Introducing the xSeries family of servers 2

Chapter 2. Linux installation 3
2.1 Knowing your hardware 3
2.2 Hardware considerations 3
2.2.1 Adaptec SCSI controller 4
2.2.2 IBM ServeRAID controller 4
2.2.3 IBM token-ring network adapters 4
2.2.4 Systems with multiple PCI buses 4
2.3 Updating the BIOS and ServeRAID controller firmware 4
2.4 Making the CD-ROM bootable 5
2.5 Basic Linux installation 5
2.6 Installation on the systems with multiple PCI buses 35
2.7 Installation with ServeRAID 36
2.7.1 Recompiling the latest kernel with a new ServeRAID driver 38
2.7.2 Recompiling the new ServeRAID driver into an existing kernel . . 42
2.7.3 Installing the Caldera OpenLinux on the ServeRAID 4 45
2.7.4 Installation with multiple PCI buses and ServeRAID 4 48
2.7.5 Installing ipsutils.rpm 52
2.7.6 The ipssend commands 55
2.7.7 Replacing a defunct drive 67
2.7.8 Replacing a defunct drive with disabled Hot Spare Rebuild 68
2.7.9 Replacing a defunct drive with a hot spare drive installed 69
2.7.10 Using the ServeRAID Manager utility 75
2.7.11 Remote management of ServeRAID adapter 78
2.8 Installing and configuring token-ring network cards 83

Chapter 3. Basic system administration 87
3.1 Log in to the system 87
3.2 Using the Window Manager 89
3.3 Getting the X-Windows terminal window 90
3.4 Accessing COAS - Caldera Open Administration System 91
3.5 Adding and removing software packages using kpackage 93
3.5.1 Uninstalling a package 94

3.5.2 Installing a package 95
3.6 Package management using RPM 97
3.7 System menu 98
3.8 Accounts 99
3.8.1 Managing accounts 101
3.8.2 Managing groups 106
3.9 Daemons (services) 109
3.10 Filesystem 110
3.10.1 Mounting an NFS volume 110
3.11 Hostname 111
3.12 Resources 111
3.13 Time 113
3.14 Peripherals menu 114
3.15 Mouse 115
3.16 Printer 116
3.16.1 Adding a new printer 117
3.16.2 Removing a printer 119
3.16.3 Edit a printer 119
3.17 Network menu 120
3.18 Ethernet interfaces 121
3.18.1 Adding a new network interface 122
3.18.2 Removing a network interface 124
3.19 Name resolution settings 124
3.19.1 Name resolution order and sources 125
3.19.2 Defining a DNS server 125
3.20 Manipulating kernel modules 127
3.20.1 Loading a new module 128
3.20.2 Unloading a new module 129
3.21 Configuring X-Windows 129
3.22 System administration using Webmin 129
3.22.1 Webmin section 132
3.22.2 Webmin system 132
3.22.3 Webmin servers 134
3.22.4 Webmin hardware 135
3.22.5 Webmin Others 137
3.23 Filesystem permissions 138

Chapter 4. General performance tools in Linux 141
4.1 General configuration hints 141
4.1.1 Services 142
4.1.2 Kernel recompilation 143
4.2 System monitoring and performance test tools 144

Chapter 5. Samba 155
5.1 What can you do with Samba? 155
5.2 Setting up the Samba server 155
5.2.1 Configuring the Samba server 157
5.2.2 Starting and stopping the Samba server 165
5.2.3 Starting Samba as startup service 166
5.2.4 Using SWAT 167
5.3 Sources of additional information 185

Chapter 6. DNS - Domain Name System 187
6.1 Installation of software 189
6.2 DNS sample configuration 190
6.3 Configuration tips 195

Chapter 7. DHCP - Dynamic Host Configuration Protocol 197
7.1 What is DHCP? 197
7.2 Why should I use DHCP? 197
7.3 Implementation on Linux 197

Chapter 8. Apache and IBM HTTP Servers 199
8.1 The IBM HTTP Server 200
8.2 Apache HTTP Server installation 201
8.3 IBM HTTP Server installation 202
8.3.1 Activating IBM HTTPD on system bootup 205
8.3.2 Setting up the Administration Server 205
8.4 General performance tips 209

Chapter 9. sendmail 213
9.1 What is sendmail? 213
9.2 What you can do with sendmail 213
9.3 Before you begin 213
9.4 Network configuration 215
9.5 Setting up the DNS configuration 216
9.5.1 Setting up the master DNS 216
9.5.2 Setting up the DNS for the first subdomain 220
9.5.3 Setting up the DNS for the second subdomain 223
9.5.4 Setting up sendmail 226
9.5.5 Configuring sendmail using Webmin 229
9.5.6 Setting up the mail client 234
9.6 Sources of additional information 237

Chapter 10. LDAP - Lightweight Directory Access Protocol 239
10.1 What is LDAP? 239
10.1.1 Directory Services 239

10.1.2 X.500 240
10.2 How can I use LDAP? 240
10.3 LDAP basics. 240
10.4 Implementation on Linux 241
10.4.1 Slapd.conf 242
10.4.2 ldap.conf. 242
10.4.3 nsswitch.conf 244
10.4.4 /etc/pam.d/login 244
10.4.5 Starting OpenLDAP 245
10.4.6 Testing authentication 245

Chapter 11. NIS - Network Information System 247
11.1 What is NIS? 247
11.2 How can I use NIS? 247
11.3 Implementation on Linux 248
11.3.1 NIS Server 249
11.3.2 NIS Client. 251
11.4 Sources of additional information 256

Chapter 12. NFS - Network File System 257
12.1 The NFS process 257
12.2 Allowing NFS access to data 260
12.3 Accessing data remotely with NFS - the command line view. 262
12.4 Allowing NFS access to data with GUI 263

Chapter 13. Packet filtering with IP Chains 265
13.1 What is packet filtering? 265
13.2 What can you do with Linux packet filtering? 265
13.3 What do you need to run packet filtering?. 266
13.4 Network configuration for a packet filtering implementation 266
13.5 How to permanently enable IP Forwarding 268
13.6 Your first IP Chains success. 272
13.7 How packets travel through a gateway 273
13.8 Using IP Chains 275
13.8.1 How to create a rule 276
13.8.2 Making the rules permanent 277
13.9 Sources of additional information 278

Chapter 14. Secure Shell 279
14.1 Installing SSH. 279
14.1.1 Installing OpenSSL. 279
14.1.2 Installing OpenSSH 280
14.2 Configuring SSH. 280
14.2.1 Host key generation 281

14.2.2 SSHD server dameon 281
14.2.3 User key generation 281
14.2.4 Configuring connections 282

Chapter 15. SNMP 285
15.1 SNMP - What is it? 285
15.2 Community strings 286
15.3 Why should I use SNMP? 288
15.4 Implementation on Linux 288
15.4.1 MRTG 291
15.4.2 Sources of additional information 295

Chapter 16. Backup and recovery 297
16.1 BRU 297
16.1.1 Installing BRU 297
16.1.2 Basic commands 299
16.1.3 Basic backup 299
16.1.4 Basic restore 299
16.1.5 Basic verification and listing commands 300
16.1.6 X Interface 301
16.1.7 The big buttons in BRU 301
16.1.8 Creating archives 302
16.1.9 Scheduling 304
16.1.10 Restoring files 305
16.1.11 Listing and verifying archives 305
16.1.12 Summary 306
16.2 Microlite BackupEDGE 306
16.2.1 Installing Microlite BackupEDGE 307
16.2.2 Initializing the tape 308
16.2.3 Your first backup 310
16.2.4 Restoring single files or directories 314
16.2.5 Master and incremental backups 316
16.2.6 Restoring master and incremental backups 319
16.2.7 Performing scheduled backups 320
16.2.8 Configuring the tape devices 323
16.2.9 Defining the devices for making backups 329
16.2.10 Microlite RecoverEDGE 332
16.2.11 More information on Microlite 342
16.3 Arkeia 342
16.3.1 Installing Arkeia 343
16.3.2 Configuring Arkeia 343
16.3.3 Interactive backup 358
16.3.4 Periodic Backup 362

16.3.5 Restoration 363
16.3.6 Advanced features of Arkeia 367

Appendix A. RAID levels 369
A.1 What is RAID? 369
A.1.1 RAID-0 370
A.1.2 RAID-1 and RAID-1E 371
A.1.3 RAID-10 372
A.1.4 RAID-5 373
A.1.5 RAID-5 enhanced 377
A.1.6 Orthogonal RAID-5 379
A.1.7 Performance 380
A.1.8 Recommendations 382
A.1.9 Summary 383

Appendix B. Working video modes for IBM Netfinity servers 385

Appendix C. Recommendations for disk partitions 387

Appendix D. Hardware issues for IBM Netfinity servers 389

Appendix E. Sample smb.conf SAMBA configuration file 391

Appendix F. Modified ifup-dhcp file 399

Appendix G. Using the additional material 401
G.1 Using the diskettes 401
G.1.1 System requirements for using the diskettes 401
G.1.2 How to use the Web material 401

Appendix H. Special notices 403

Appendix I. Related publications 407
I.1 IBM Redbooks 407
I.2 IBM Redbooks collections 407
I.3 Other resources 407
I.4 Referenced Web sites 408

How to get IBM Redbooks 411
IBM Redbooks fax order form 412

Index 413

IBM Redbooks review 419

Preface

Caldera OpenLinux eServer 2.3 is the latest edition of the popular Caldera Linux distribution. OpenLinux eServer 2.3 is a fully tested, proven, stable and supported Linux distribution for corporate and home use. OpenLinux delivers reliability, ease of installation and administration, high performance, security, and robust applications. This version of OpenLinux includes the latest version of LIZARD (Linux Wizard), the powerful but simple-to-use graphical installation program. LIZARD automatically detects all supported hardware and helps you to get OpenLinux up and running quickly.

This redbook will help you install and configure OpenLinux on your IBM @server xSeries and Netfinity servers; furthermore it will help you understand, install and configure a wide range of different services, such as Samba, Apache and LDAP, among others.

Do I need to be a Linux expert to use this guide? The answer to that question is “No”. This redbook is aimed at beginners and intermediate Linux users and for all Windows NT users who are used to the safe and convenient graphical user interface.

The team that wrote this redbook

Figure 1. The team (left to right) Credle, Holzknecht, 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 in the IBM Software Group. His most recent publication was *Small Business Suite for Linux Reviewer's Guide*. Before joining the IBM Software Group, he worked at the International Technical Support Organization center in Raleigh, 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 Electronics 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.

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.

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 intergration. He joined SuSE in May of 2000 after graduating from the University of Derby with a diploma in computer science.

Ivo Gomilsek is an IT Specialist for Storage Area Neworks 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, Implementing Vinca Solutions on IBM Netfinity Servers*, and the first edition of *Netfinity and Linux Integration Guide*. He also provides Level 2 support for IBM Netfinity servers and high availability solutions for IBM Netfinity servers and Linux. Ivo has been employed at IBM for four years.

Georg Holzknecht 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 original authors, Lenz Grimmer and Joe Kaplenk, for their contribution to the first edition of this redbook, which was titled *Caldera OpenLinux and Netfinity Server Integration Guide*, SG24-5861-00.

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

Brad Dew, Business Development Manager
Caldera Systems, Inc.

Thanks to the following IBM employees:

Egan Ford, Advanced Technical Support

Karl Schultz, Netfinity ServerProven

Megan Chilton, Marketing Communications

Bo Brun, PC Institute

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 419 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

IIBM 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 Caldera Systems, Inc.

Caldera Systems, Inc. was founded in 1994 by Ransom Love and Bryan Sparks. In 1998, Caldera Systems, Inc. was created to develop Linux-based business solutions. Based in Orem, UT, Caldera Systems, Inc. is a leader in providing Linux-based business solutions through its award-winning OpenLinux line of products and services.

OpenLinux is a full-featured, 32-bit, "Linux for Business" solution that expands Internet/intranet, network and desktop capabilities using the open source technologies of the Linux kernel. Through OpenLinux solutions, the total cost and development for small-to-medium size businesses is greatly reduced. Caldera Systems provides free 30-day phone support plus e-mail and fee-based 24 x 7 support to VARs and end users, including Linux systems administration training and certification courses through Caldera Systems' OpenLearning program worldwide. Caldera Systems is the first Linux company to create the Linux VAR channel with over 1000 resellers worldwide and is one of IBM's chosen distribution partners for Linux.

The IBM @server xSeries Brand team works closely with Caldera Systems and our other distribution partners to fully test and certify that xSeries and Netfinity servers are ready to perform with Linux.

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. Linux installation

This chapter describes in detail the basic installation necessary to install and run Caldera OpenLinux eServer 2.3 on IBM @server xSeries and Netfinity servers. We also show what steps you need to carry out before the installation to avoid problems during the installation.

2.1 Knowing your hardware

Before you install Caldera OpenLinux, you should be familiar with the hardware in the computer. You at least need information about the following components in your computer:

- Hard drives - interface type (SCSI or IDE) and size
- CD-ROM - interface type (SCSI or IDE) and the manufacturer
- SCSI adapter - manufacturer and model number
- Display adapter - manufacturer and model number
- Mouse - mouse type and connector type
- Network card - manufacturer and model number
- RAM - the amount of RAM in your system
- Monitor - manufacturer and model number

For the monitor in particular, it is good to know all technical specifications.

For IBM Netfinity servers and other IBM products, including monitors and SCSI adapters, the ultimate source for all information is:

```
ftp://ftp.pc.ibm.com/pcicrse/psref
```

Here you will find the PSREF (Personal Systems Reference) sheets for all IBM PC products, both current and withdrawn. You can also find a lot of useful information at the following Internet resources:

```
http://www.pc.ibm.com/support
http://www.pc.ibm.com/us/netfinity/tech_library.html
```

2.2 Hardware considerations

In this section we describe how to handle hardware found in IBM xSeries and Netfinity servers, when installing Caldera OpenLinux.

2.2.1 Adaptec SCSI controller

All versions of the Adaptec SCSI controllers used in xSeries and Netfinity servers are supported in kernels 2.2.14 and above, which are used in Caldera OpenLinux eServer 2.3.

2.2.2 IBM ServeRAID controller

The IBM ServeRAID controller up to Version 3 used in xSeries and Netfinity servers are supported in Caldera OpenLinux eServer 2.3. For support for the ServeRAID controller Version 4 refer to 2.7, “Installation with ServeRAID” on page 36 for more information.

2.2.3 IBM token-ring network adapters

IBM ISA, PCMCIA and PCI token-ring network adapters are supported in kernels 2.2.14 and above, which are used in Caldera OpenLinux eServer. The token-ring PCI LAN Streamers adapters are not supported in the kernel 2.2.14 used in Caldera OpenLinux eServer 2.3. To use these adapters you need to use the latest available kernel from Release 2.2.x and follow the instructions in 2.7.1, “Recompiling the latest kernel with a new ServeRAID driver” on page 38 and 2.8, “Installing and configuring token-ring network cards” on page 83 for successful installation.

2.2.4 Systems with multiple PCI buses

If you have a system with multiple PCI buses you will need to obtain a special installation disk and some updated packages from Caldera Systems FTP server. Follow the instructions in 2.6, “Installation on the systems with multiple PCI buses” on page 35 for installation on such systems.

2.3 Updating the BIOS and ServeRAID controller firmware

Before starting the installation it is important to have the latest level of microcode for all your hardware components. You can download the latest BIOS, diagnostic updates and ServeRAID controller firmware for your xSeries or Netfinity server from:

```
http://www.pc.ibm.com/support
```

For the networking products, you will find all the latest available code at:

```
http://www.networking.ibm.com
```

Note

Always update the BIOS of the xSeries and Netfinity server to the latest level, and update all adapters with the latest firmware before installing.

2.4 Making the CD-ROM bootable

If you want to use the CD-ROM for booting the system to start the installation make sure the CD-ROM is in the boot sequence before any hard drive devices. You can modify this in the IBM Netfinity server BIOS. To do this follow 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 > Startup Sequence**.
4. Make sure that your CD-ROM is in the boot sequence before the first hard disk device.
5. Press Esc until you come to the setup utility main window and select **Save Settings.**
6. Press Enter to confirm saving the current settings.
7. 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.

2.5 Basic Linux installation

In this section we will describe how to install the Caldera OpenLinux eServer 2.3 on xSeries and Netfinity servers. To successfully complete the installation follow these steps:

Note

With this book, you received the updated CD for use with the latest IBM xSeries and Netfinity server. If you are having troubles installing with the original CDs you should try to install with this updated CD. The CD is also available for download from the Caldera Systems FTP site: `ftp://ftp.calderasystems.com`. You can find it under **`updates`** for the eServer product.

1. Insert Caldera OpenLinux eServer 2.3 into the CD-ROM drive of your IBM xSeries and Netfinity server. When the CD is started, you will see a window similar to Figure 2.

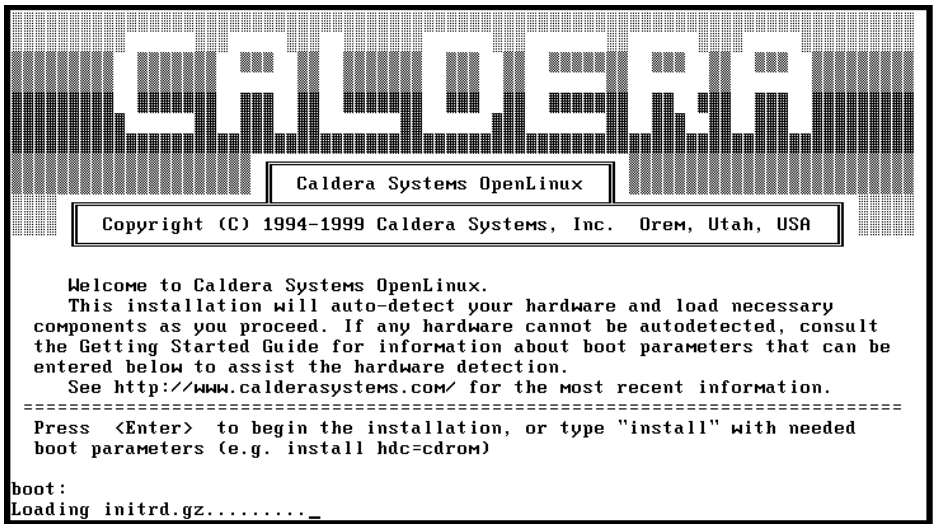

Figure 2. Caldera OpenLinux bootup

After the initial RAM disk and kernel are loaded you will see a window similar to Figure 3.

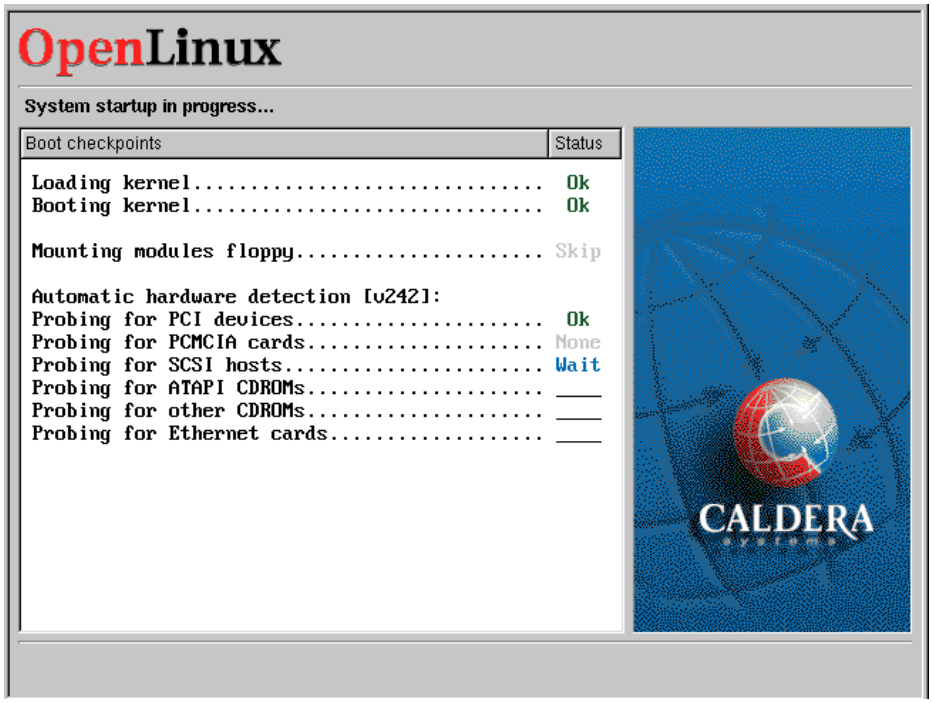

Figure 3. Hardware detection

When all the hardware is detected you will see a window similar to Figure 4.

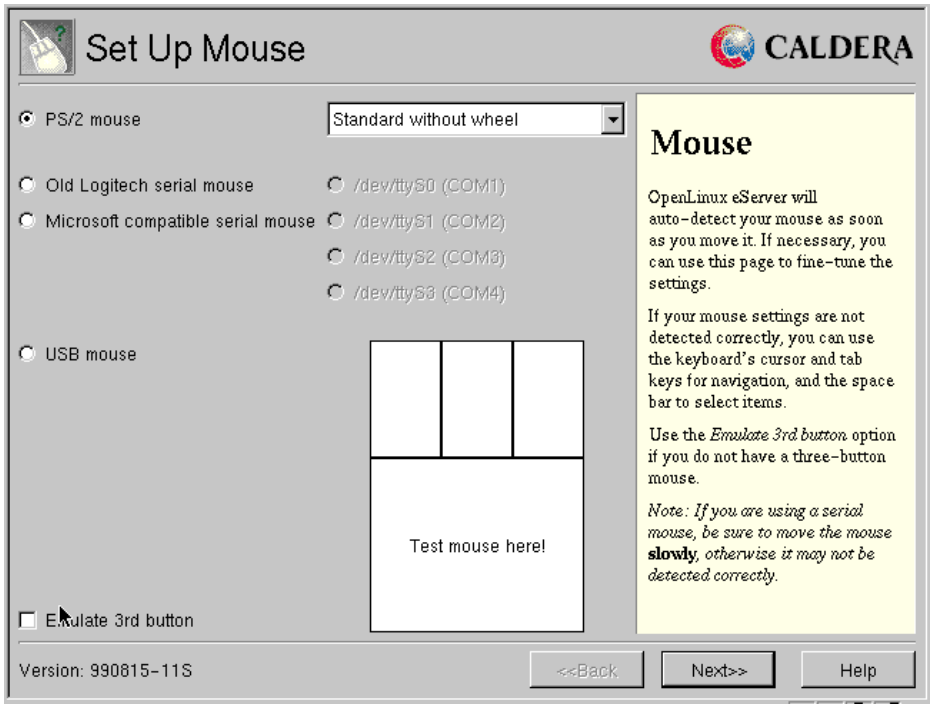

Figure 4. Mouse setup

2. Here you select your mouse type and settings. When you are finished, click **Next** to continue. You will see a window similar to Figure 5.

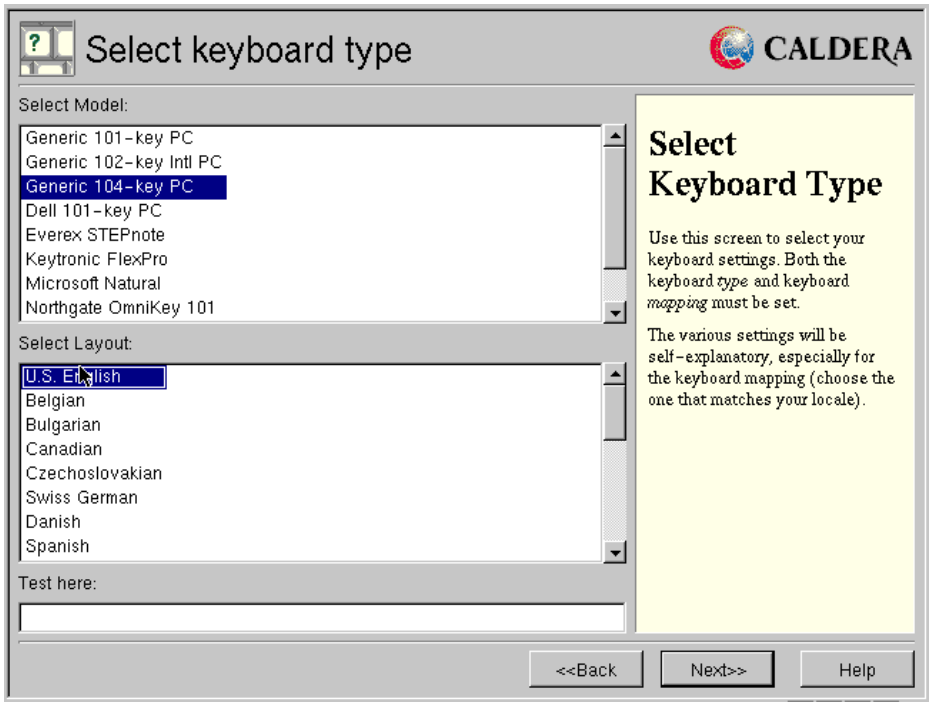

Figure 5. Keyboard setup

3. In this window you select the keyboard type and layout. Click **Next** to continue. You will see a window similar to Figure 6.

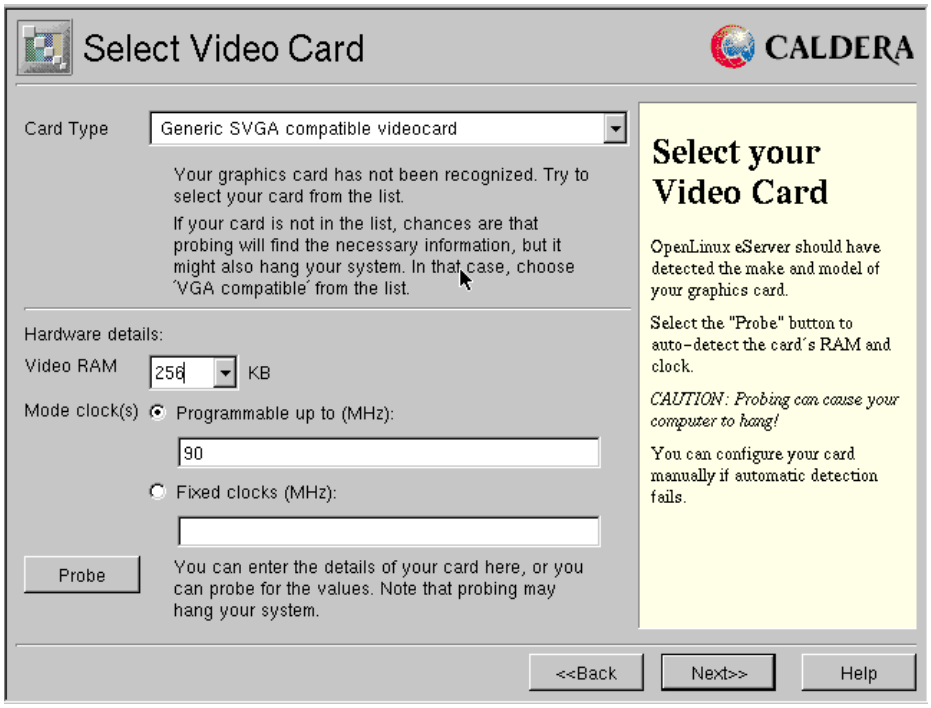

Figure 6. Video card setup

4. Here you select the XFree server you will use for your video card. Before selecting your XFree video server check Appendix B, “Working video modes for IBM Netfinity servers” on page 385. When defining the card type, the install program automatically picks up the XFree server that supports that card. The best way is to select, from the Card Type scroll box, the XFree server that supports the card you are using. After you have selected the XFree server you want, you need to specify the amount of Video RAM and also the Mode clock. If you do not know the exact values, you can click **Probe** to search for the details of the graphic card.

Note

Before successfully probing for the card, you need to specify the correct XFree server for your card. If you do not know which server to use you should use the **Generic SVGA compatible videocard** option.

If you click **Probe**, you will see a window similar to Figure 7. If you select the VGA XFree server, the installation procedures will continue with step 11.

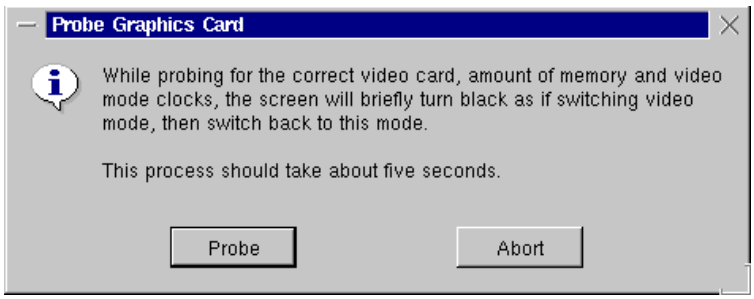

Figure 7. Before probing for the video card

5. Click **Probe** to start probing for the card. The window will briefly turn black. When probing is finished, you will see a window similar to Figure 8.

Figure 8. After successfully probing for the card

6. If for some reason the video card cannot be recognized correctly, you need to type in the settings manually. If you do not know the correct settings, select the VGA XFree server and you can fine-tune XFree after the installation is completed. Click **OK** to return to the video card setup window. Click **Next** to continue. You will see a window similar to Figure 9.

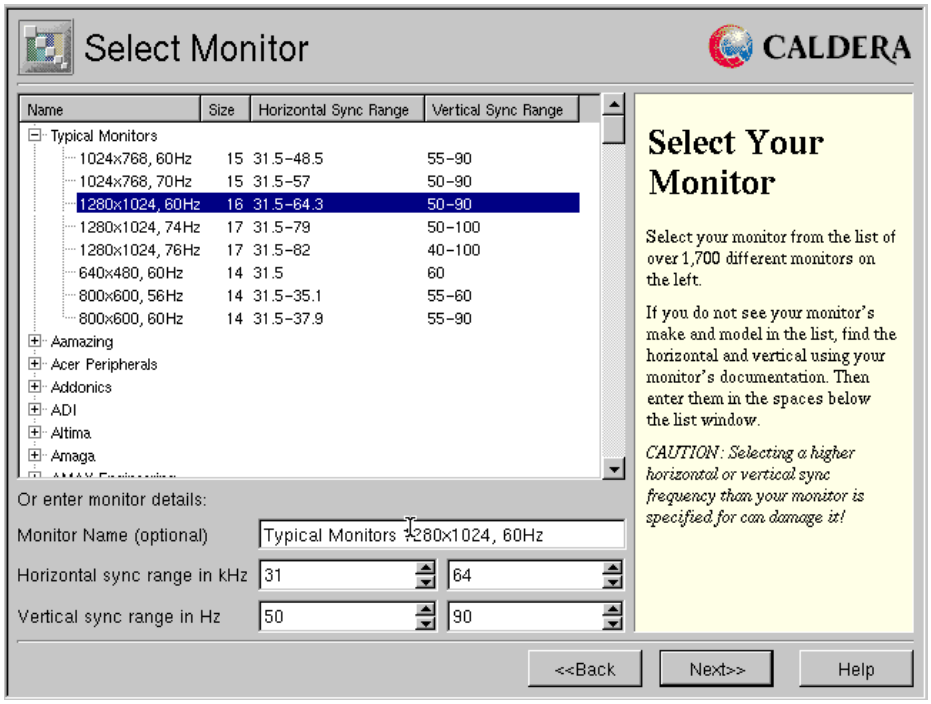

Figure 9. Select Monitor

7. Select the monitor type you are using. If there is no exact match for your monitor model, select one of the Typical Monitors. Use the one nearest to the one you are using. Click **Next** to continue, and you will see a window similar to Figure 10.

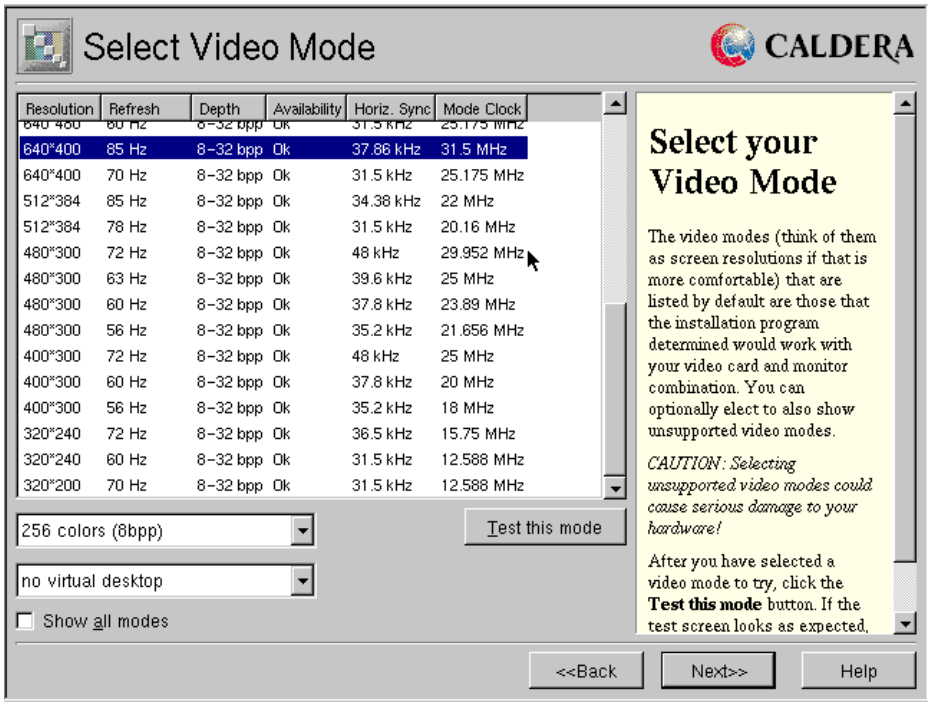

Figure 10. Select Video Mode

8. In this window you select the video mode, color depth and the virtual desktop size you prefer. You can find working modes for IBM Netfinity servers in Appendix B, “Working video modes for IBM Netfinity servers” on page 385. Depending on the amount of memory on the video card, color depth and resolution, you can have a virtual desktop. This is a feature of the XFree servers. By default only the modes that are compatible with the selected video card and monitor are displayed. If you are an experienced user you can check the **Show all modes** option to see all available video modes. You can test the selected video mode by clicking **Test this mode.** You will see a window similar to Figure 11.

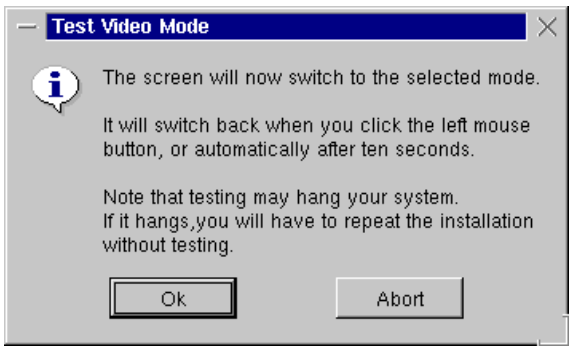

Figure 11. Test Video Mode

9. Click **OK** to start the test. If you have selected a working mode for your hardware, you will get a clear demo image and a message how to get back to the installation process. This image will disappear after 10 seconds if you do nothing. After you have successfully tested your video mode you will be back to the window similar to Figure 12. Click **Next** to continue the installation. You will see a window similar to Figure 12.

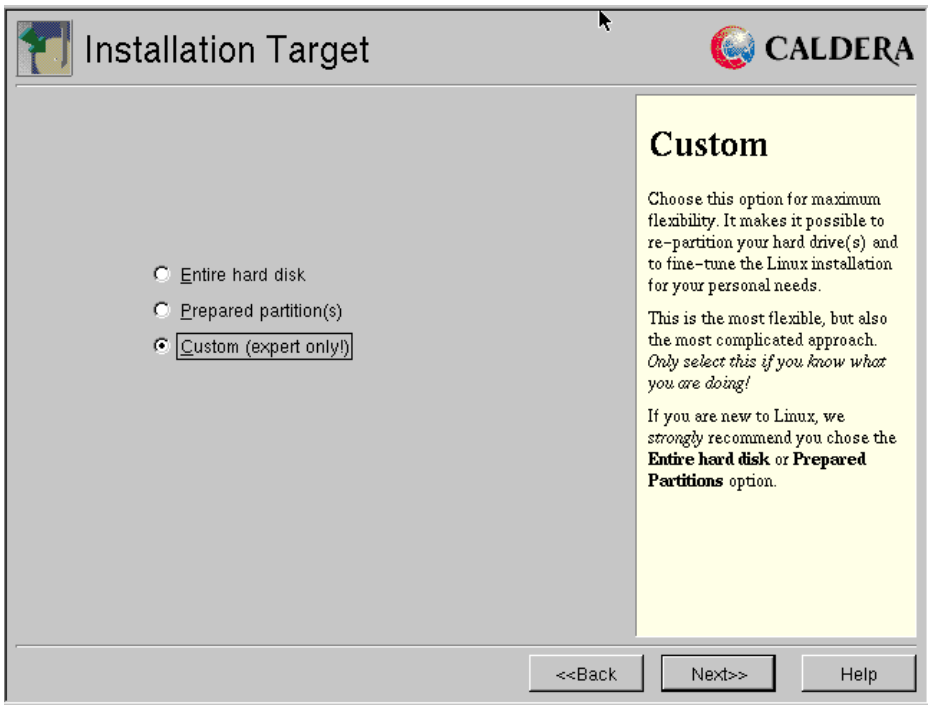

Figure 12. Choosing installation target

10. At this point you can select the type of partitioning you want to use. You have three options:

    - **Entire hard disk**: this option is initially selected if no previous Linux installation is found on the system. It means that the whole disk will be used for the Caldera OpenLinux installation.
    - **Prepared partition(s)**: this option is enabled only when a previous Linux installation is found on the hard disk. If you choose this option, the installation process will use the existing partitioning.
    - **Custom (expert only!)**: use this option if you want to create your own partition scheme.

    In our case we selected **Custom**. Click **Next** to continue the installation. You will see a window similar to Figure 13.

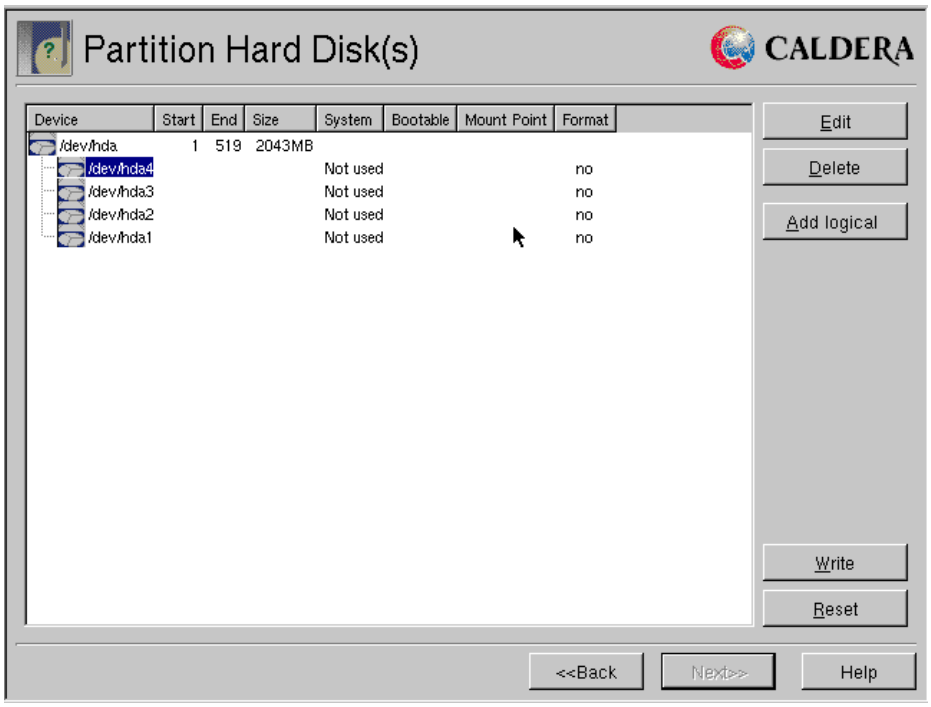

Figure 13. Partition Hard Disk(s)

11. Here you can create, delete, or modify partitions. To create a new partition select an entry that indicates it is “Not used” and click **Edit**, and you will see a window similar to Figure 14.

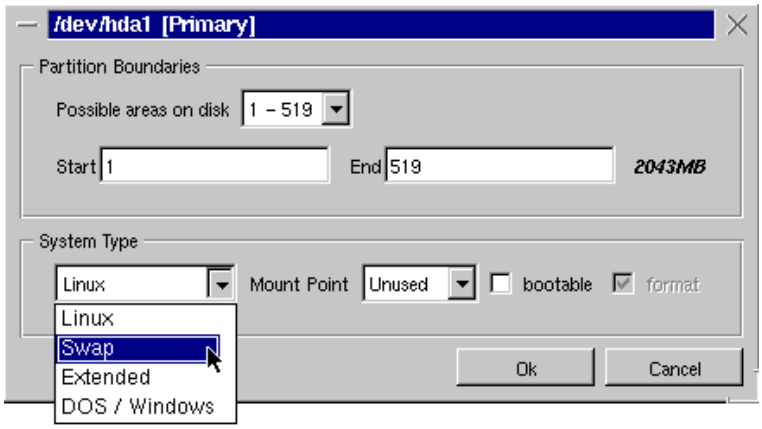

Figure 14. Creating new partition

As you can see in Figure 14, you can create four types of partitions:

- Linux
- Swap
- Extended
- DOS/Windows

To successfully install Caldera OpenLinux, you need to create at least one Linux and one swap partition. But we recommend that you create more Linux partitions for the installation files:

- One small partition for “/boot” mounting point, around 20 MB for kernel images
- Partition for “/” mounting point (ROOT), at least 1 GB
- Partition for “/home” mounting point; the size depends on your needs
- Partition for “/usr” mounting point; the size depends on your needs
- Partition for “/var” mounting point; the size depends on your needs
- Partition for “/opt” mounting point; the size depends on your needs

For the Linux type of partitions, you can select the mounting point by clicking the arrow to the right of the Mount Point field, and you will see a window similar to Figure 15.

Note

Refer to Appendix C, “Recommendations for disk partitions” on page 387 for our recommendations regarding disk partitions.

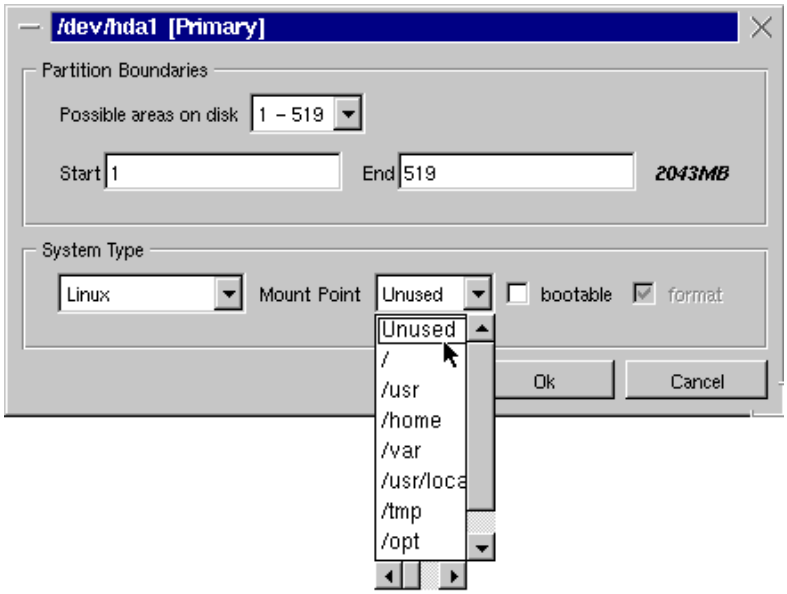

Figure 15. Select mounting point

If the mount point you need is not available, you can easily type in whatever you need. That is the case for the “/boot” mount point. If you want the partition you are creating to be bootable, select the **bootable** checkbox.

12. When you have selected the mount point and defined the size of the partition, click **OK** to save the selections you made. You will see a window similar to Figure 16.

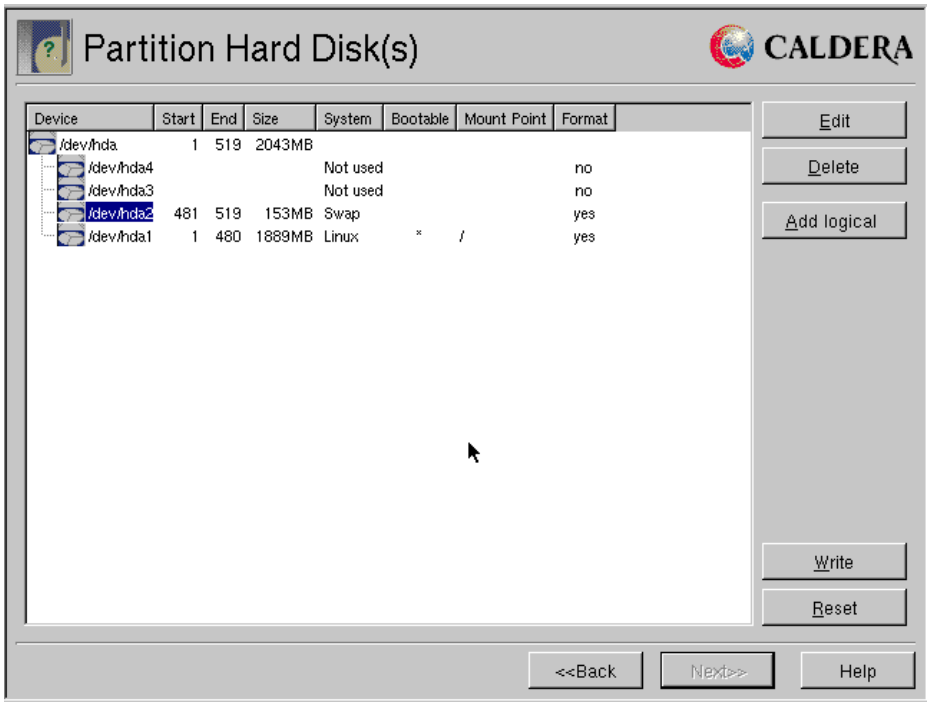

Figure 16. After defining partitions

If you need to create more than four partitions on one physical disk, this can be done by creating an extended partition with more logical drives in it. To create an extended partition you need to use the Extended type as the partition type.

13. To create an extended partition, select an unused partition entry and click **Edit**, and you will see a window similar to Figure 14 on page 16.
14. Select **Extended** as the partition type and define the partition size. Click **OK** to continue, and you will see a window similar to Figure 17.

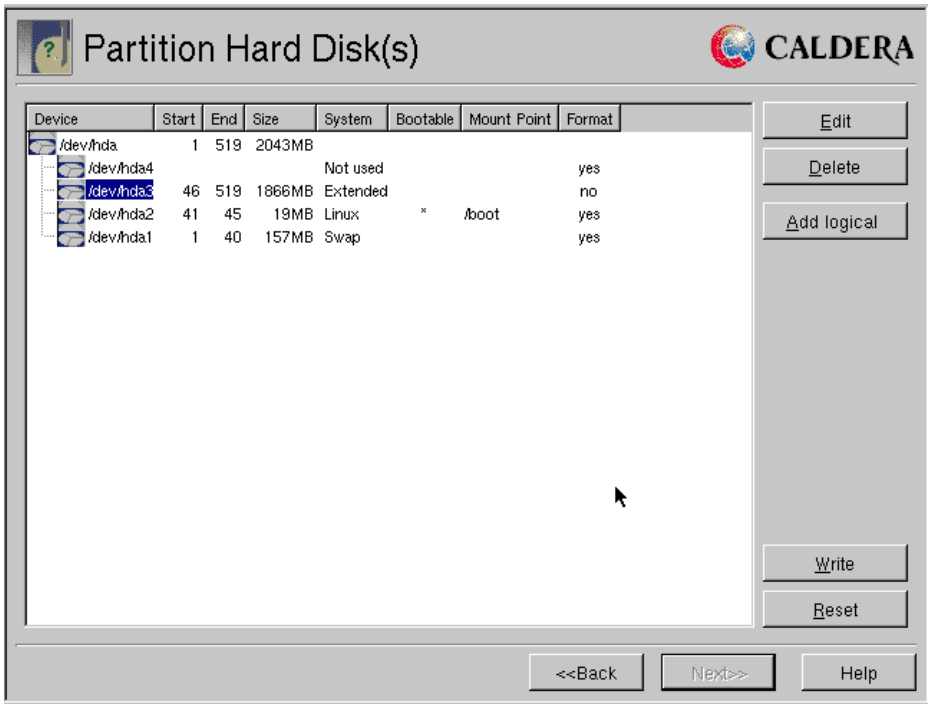

Figure 17. After creating an extended partition

As you can see in our example, there is an extended partition under device “/dev/hda3”.

15. To create logical drives inside an extended partition, select the extended partition and click **Add logical.** You will see a window similar to Figure 18.

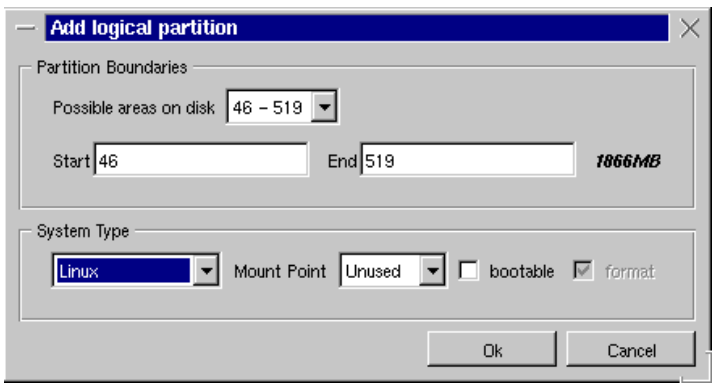

Figure 18. Add a logical partition

16. Here you specify the size and mount point of the logical partition. Then you click **OK** to continue. You will see a window similar to Figure 19.

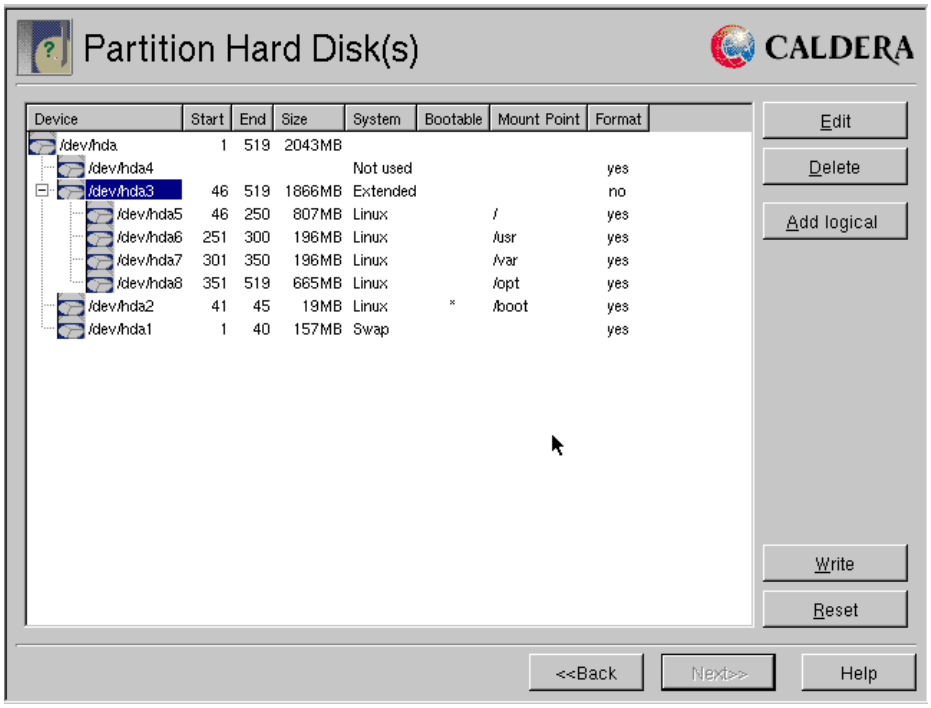

Figure 19. After creating partitions

17. The partitions you have just created are not yet written to the disk, so you can still make changes. You can click the **Reset** button to restore the original configuration. When you are finished with the changes, click **Write** to commit the changes to the disk drive. Click **Next** to continue. If you forget to assign "/" (the root partition), you will see a window similar to Figure 20, and you proceed on step 18. Otherwise, you will see a window similar to Figure 21 on page 22 and you continue at step 19.

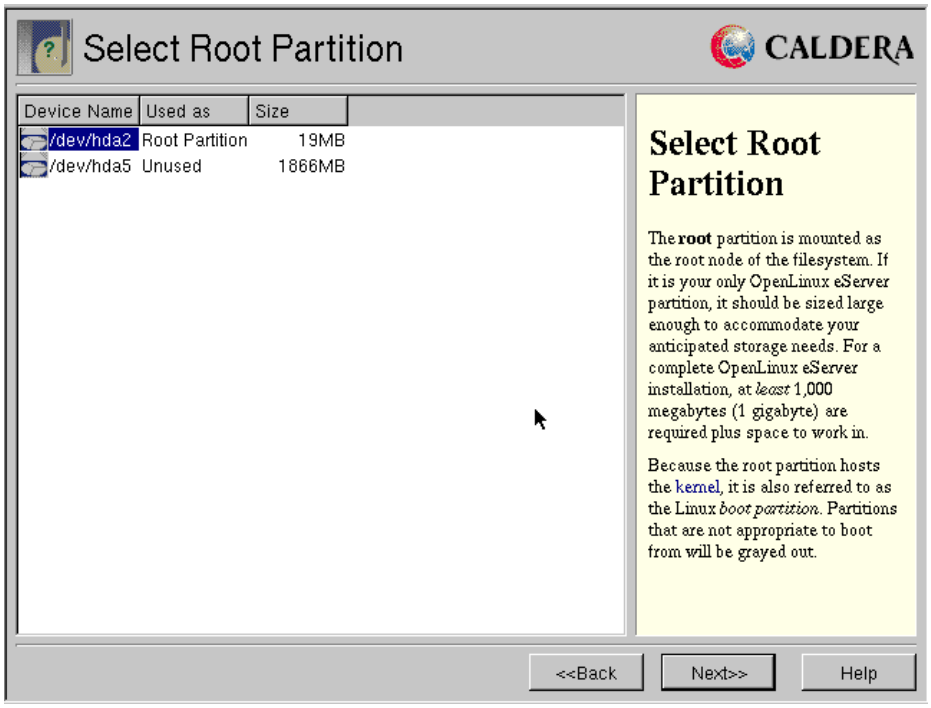

Figure 20. Select Root Partition

18. Select the partition that will be used for the root mounting point "/". Click **Next** to continue. You will see a window similar to Figure 21.

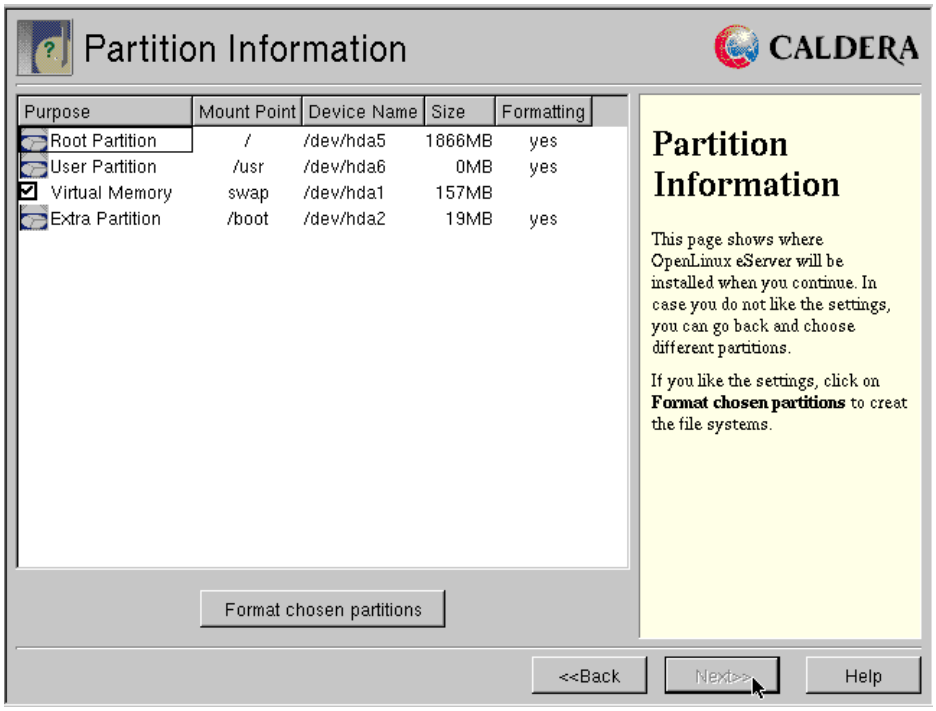

Figure 21. Select partitions to format

19. Before you can start installing the file system it has to be created on the target partitions. Do this by formatting the partitions. In this window you can select which partitions you want to format. Some partitions (for example "/" - the root partition) will be formatted by default and you cannot choose not to format them. To start formatting, click **Format chosen partitions.** You will see a window similar to Figure 22.

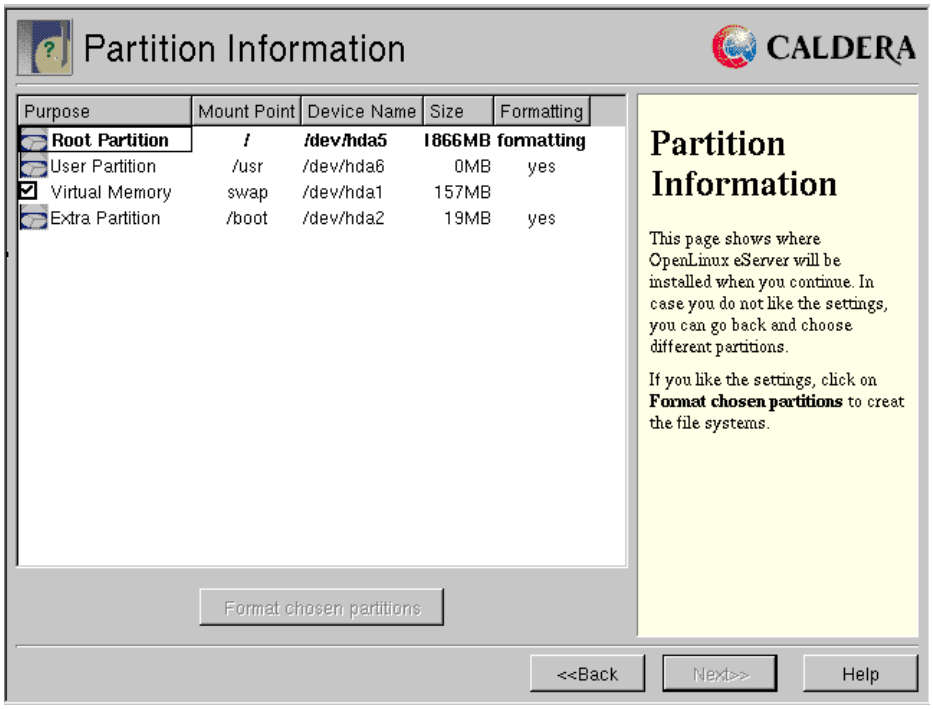

Figure 22. Formatting the partitions

After the partitions are formatted, you will see a window similar to Figure 23.

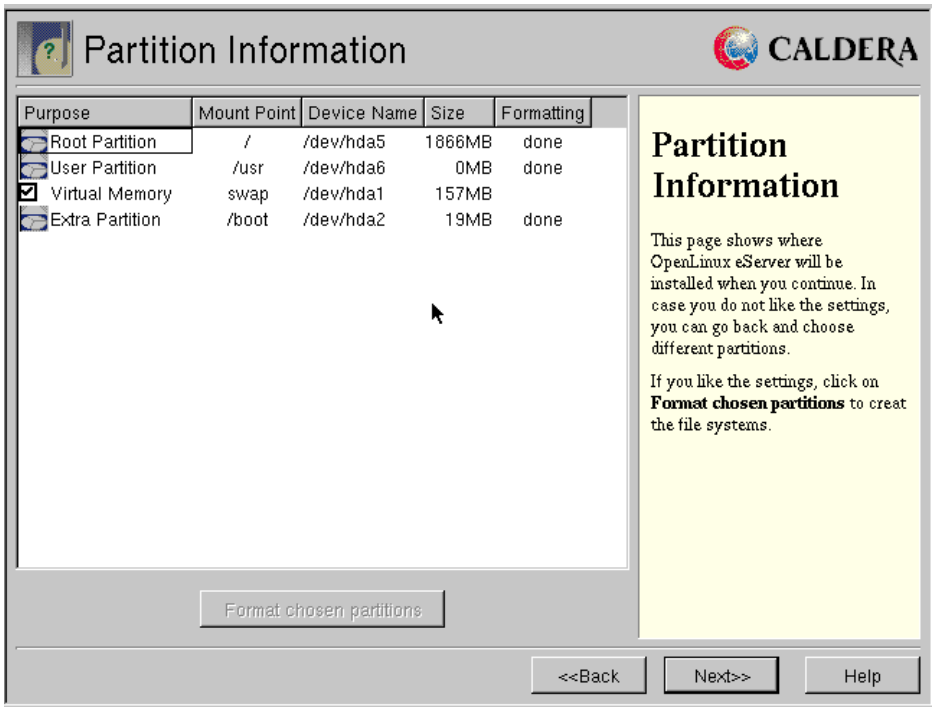

Figure 23. After formatting

20. Click **Next** to proceed with the installation. You will see a window similar to Figure 24.

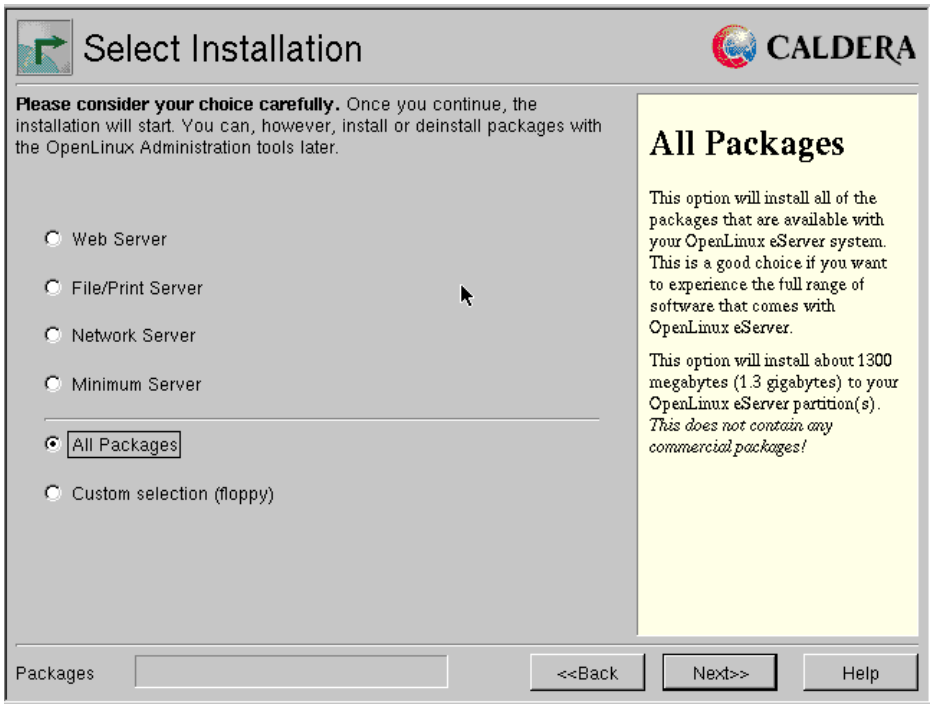

Figure 24. Select the type of installation

21. Here you select the installation type, and define which packages will be installed. If some options are grayed out, it means you do not have enough free space available for the installation. You can click the **Back** button to go back and change the partitions to meet your needs. After selecting the type of installation, click **Next** to continue. The packages will start installing in the background; Linux is a real multitasking operating system. You can see the progress in the Packages field. You will see a window similar to Figure 25.

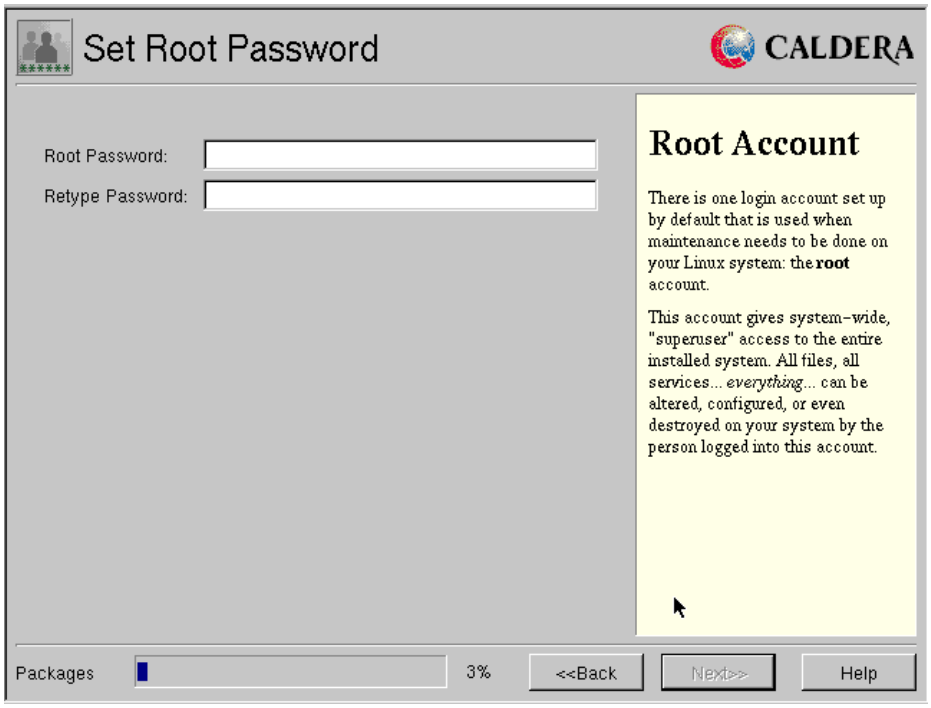

Figure 25. Defining root password

22. In this window you define your root account password. The root user will manage your system. After you entered the password, click **Next** to continue. You will see a window similar to Figure 26.

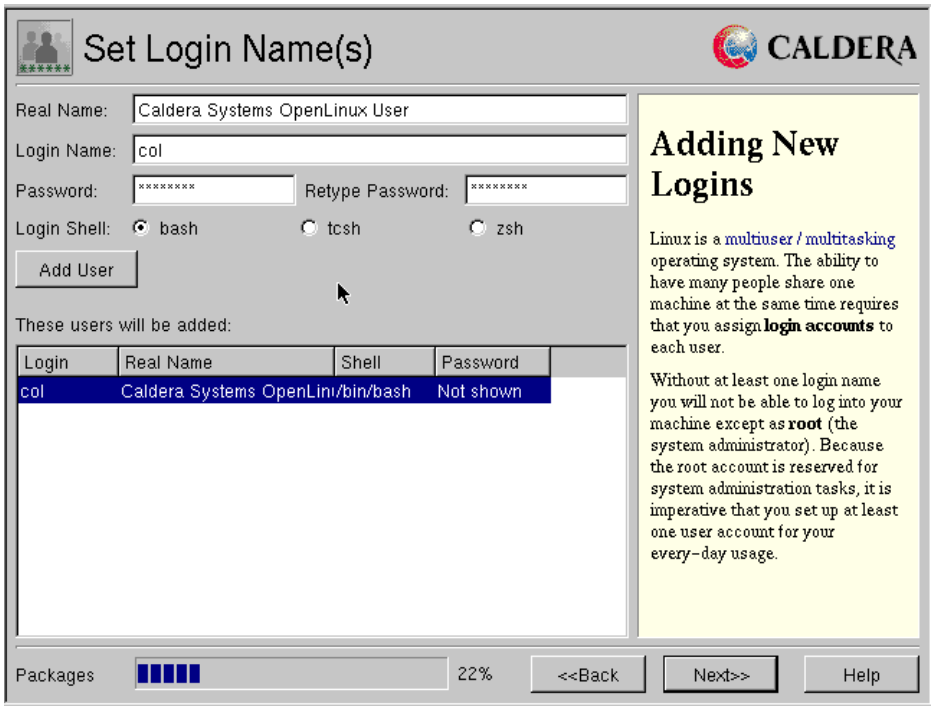

Figure 26. Defining users

23. You need to create at least one user to continue with the installation. To add a new user, fill in all fields and click **Add User.** You will see a window similar to Figure 27.

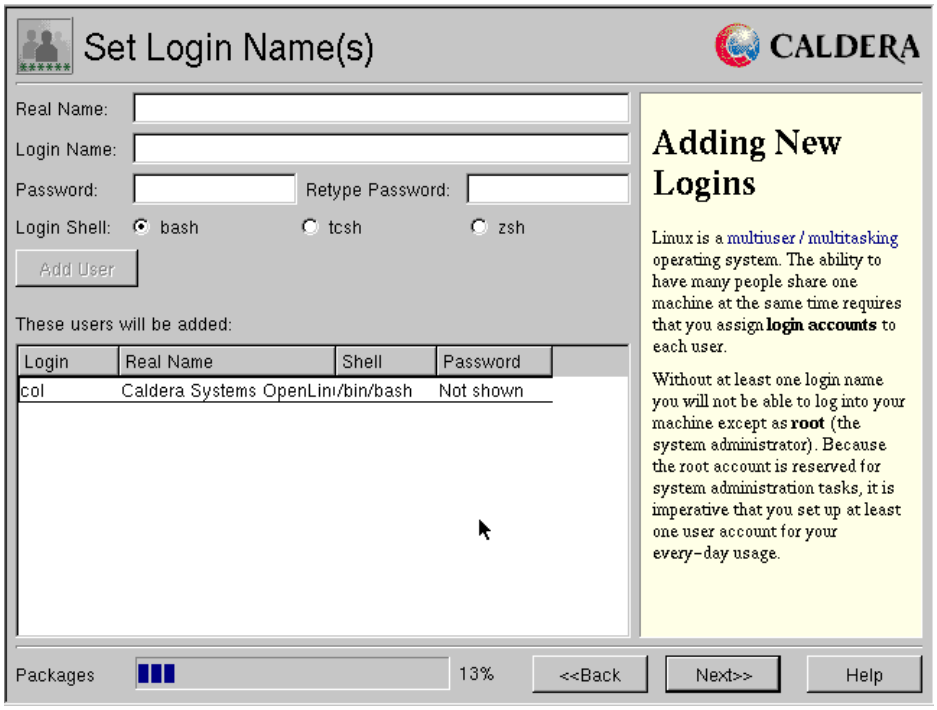

Figure 27. Adding new user

24. You can see which users have been created already. To continue with the installation, click **Next,** and you will see a window similar to Figure 28.

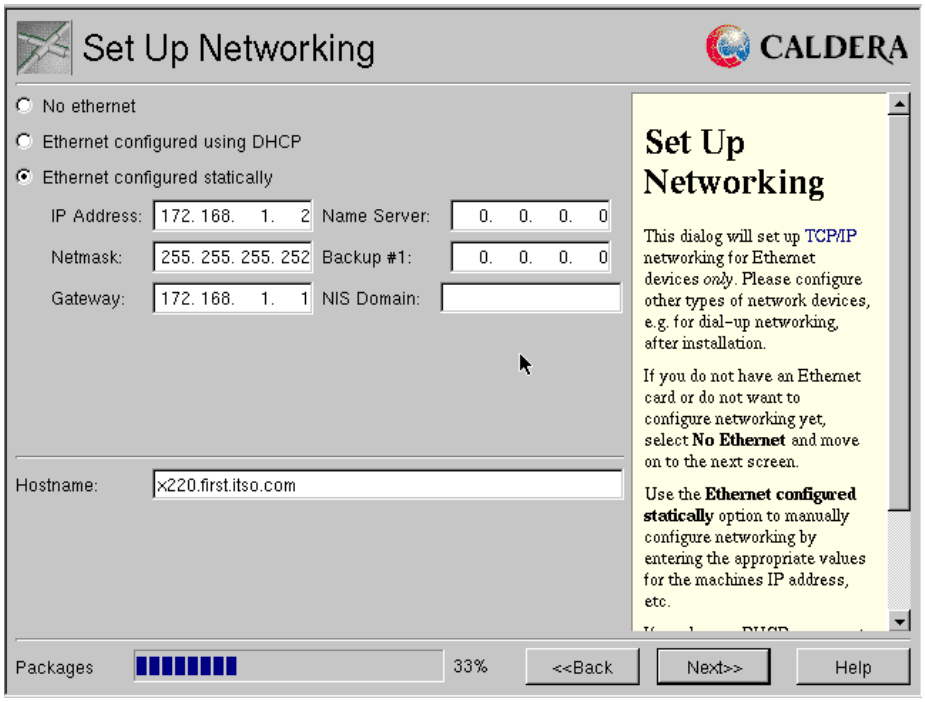

Figure 28. Set up networking

25. If your network card was recognized by the kernel, you can set it up here. The setup can be done only for Ethernet cards. You have three options:

    a. **No Ethernet**: use this option if you do not have a network card or the card is not supported, as for example token-ring. You can still configure token-ring cards later.

    b. **Ethernet configured using DHCP**: use this option when you use a DHCP server to define the IP addresses of the clients.

    c. **Ethernet configured statically**: use this option when you have a static IP address.

    On this page you also define the hostname of the computer you are installing on. All the settings can later be modified using Caldera Open Administration System (COAS). When you have typed in all needed data, click **Next** to continue. You will see a window similar to Figure 29.

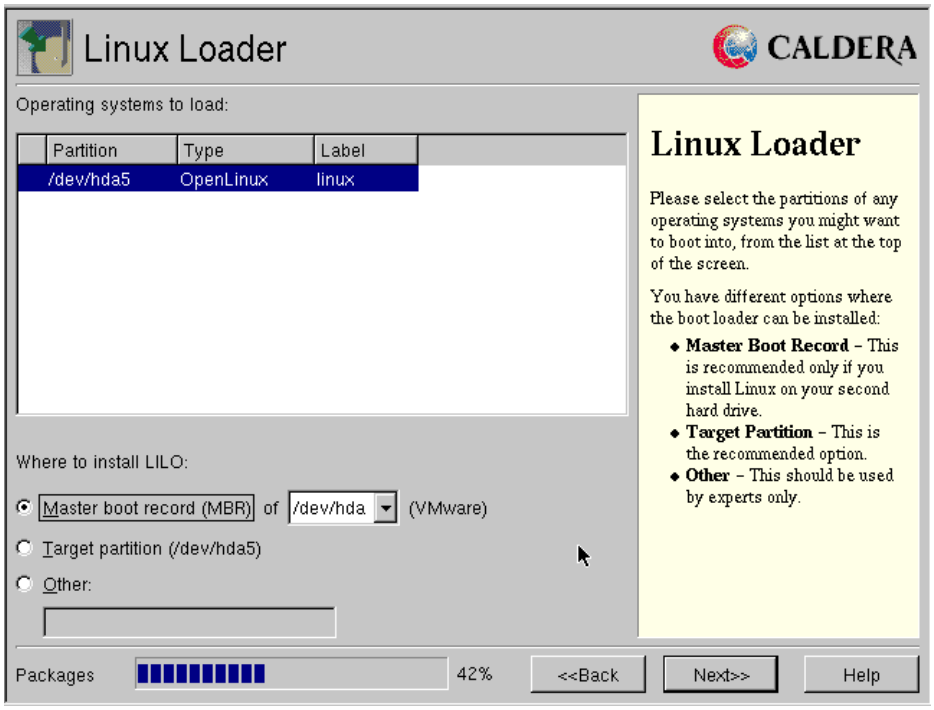

Figure 29. Configuring the Linux Loader

26. The Linux operating system is loaded by the Linux Loader (LILO). In this window you specify where the LILO will be installed. You have three options:

    a. **Master boot record (MBR)**: the LILO will reside on the first bootable physical hard drive. You use this option when Linux is the only operating system on the system or if you want to use the LILO for booting other operating systems as well.

    b. **Target partition:** in this case the LILO is set up in the boot record of the installation partition, the one that was defined as bootable. You choose this option when you are using another operating system loader (Boot Manager, System Commander, etc.), and you will be booting Linux with it.

    c. **Other:** this is the option for experienced users only.

    After you have selected your option, click **Next** to continue. You will see a window similar to Figure 30.

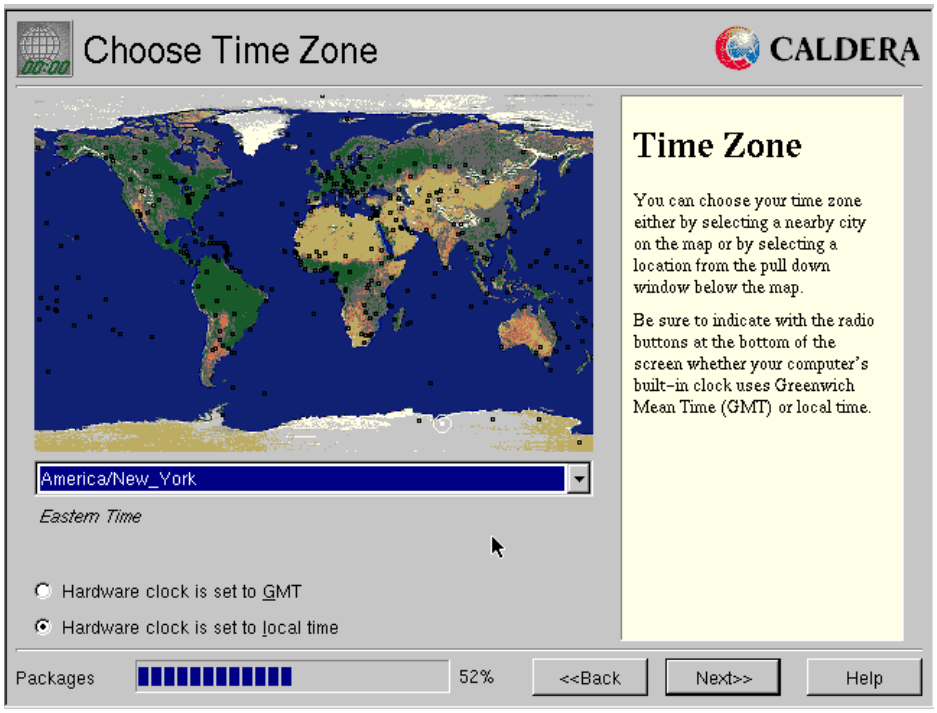

Figure 30. Time zone setup

27. Select your time zone and define how the hardware clock should be set. Click **Next** to continue, and you will see a window similar to Figure 31.

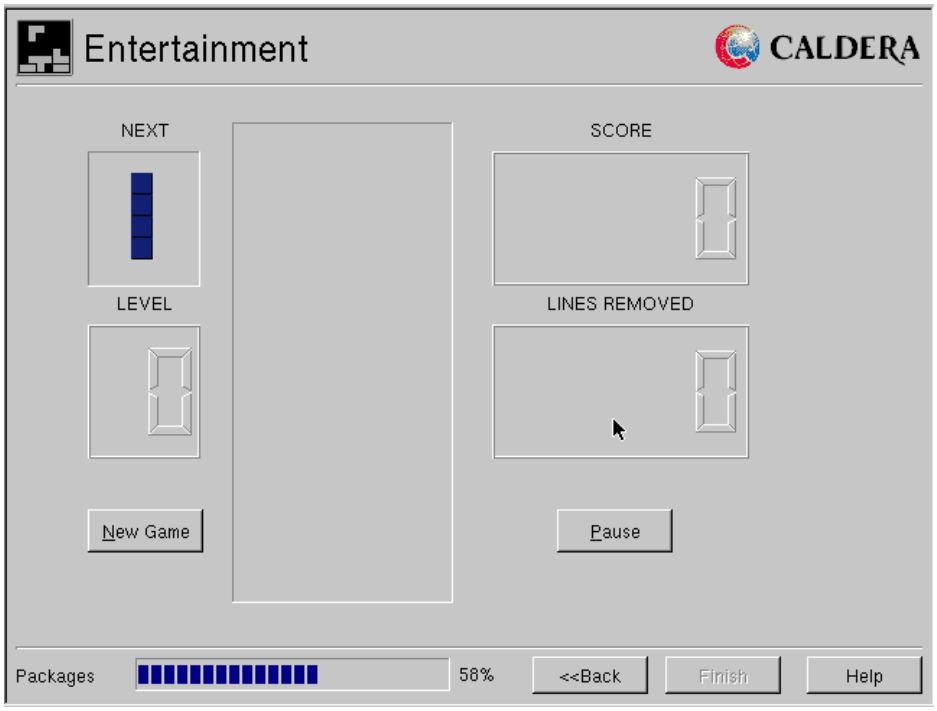

Figure 31. Entertainment during installation

28. While waiting for all the packages to be installed, you can entertain yourself by playing Tetris - the ultimate game. The installation is completed when the Finish button becomes selectable. Click **Finish** to continue, and you will see a window similar to Figure 32.

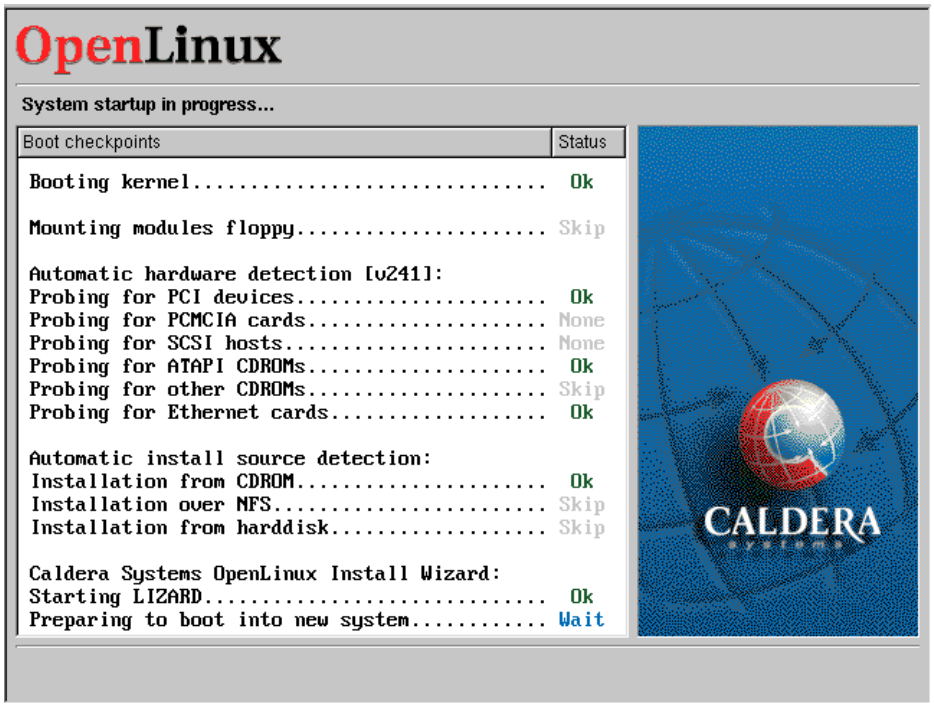

Figure 32. Booting after installation

When the system boots up, you will see a window similar to Figure 33.

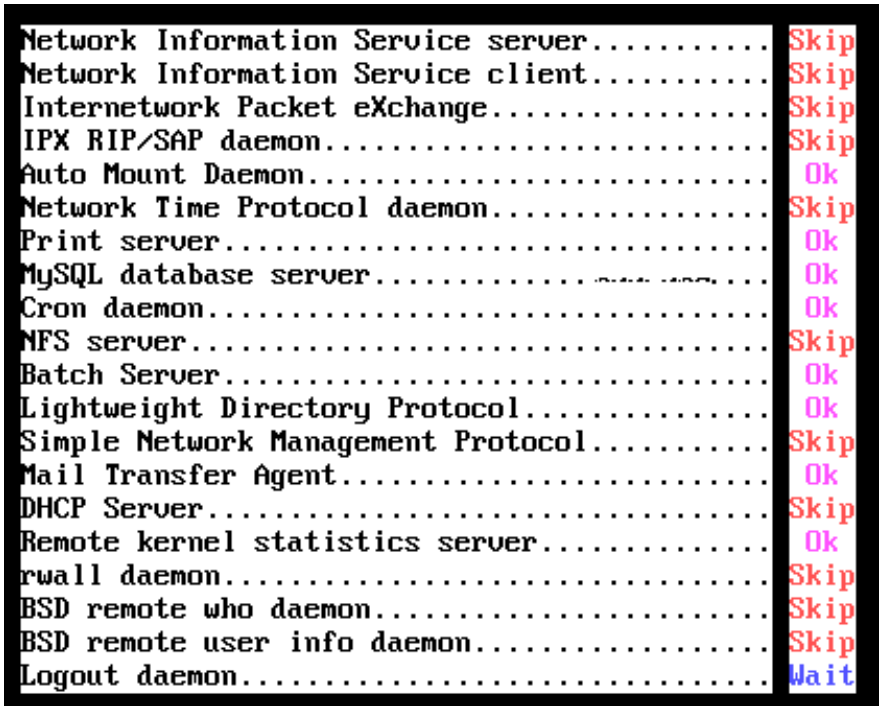

Figure 33. Starting up the operating system

When all the services are started, you will see a logon window similar to Figure 34.

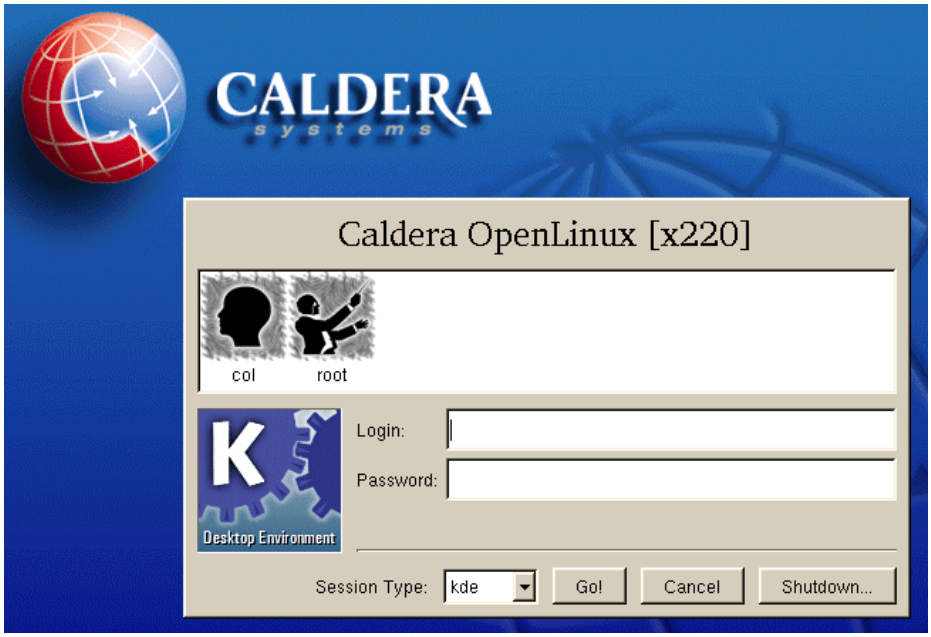

Figure 34. Logon window

Congratulations, your Caldera OpenLinux system is now ready to use!

2.6 Installation on the systems with multiple PCI buses

If you have the system with multiple PCI buses or you are using the ServeRAID adapter with BIOS Version 4.00.06 or earlier, you need to follow these instructions to successfully install Caldera OpenLinux eServer 2.3:

1. Obtain the updated installation CD and special installation disk for IBM systems with multiple PCI buses from Caldera Systems FTP site:

   `ftp://ftp.calderasystems.com/pub/eServer/updates/2.3/extras/IBM/launch/floppy/install-2.2.14-244-syslinux.img`

   You can build the floppy from another Linux system with the command:

   `dd if=install-2.2.14-244-syslinux.img of=/dev/fd0`

   The ISO image of the updated installation CD can be found in:

   `ftp://ftp.calderasystems.com/pub/eServer/updates/2.3/extras/IBM/iso-images`

2. Insert the updated installation CD into the server, and boot your server with this installation disk. Continue installation following the instructions in

2.5, “Basic Linux installation” on page 5. When the installation finishes and the system is started switch to the text console using the CTRL+ALT+F1 key combination and log on as a root.

3. From the Caldera Systems FTP site you need to obtain the following packages:

   ftp://ftp.calderasystems.com/pub/eServer/updates/2.3/extras/IBM/RPMS

   am-utils-6.0-7S.i386.rpm
   hwprobe-991112-2S_NF.i386.rpm
   linux-kernel-binary-2.2.14-5S_NF.i386.rpm
   linux-kernel-doc-2.2.14-5S_NF.i386.rpm
   linux-kernel-include-2.2.14-5S_NF.i386.rpm
   linux-source-alpha-2.2.14-5S_NF.i386.rpm
   linux-source-arm-2.2.14-5S_NF.i386.rpm
   linux-source-common-2.2.14-5S_NF.i386.rpm
   linux-source-i386-2.2.14-5S_NF.i386.rpm
   linux-source-m68k-2.2.14-5S_NF.i386.rpm
   linux-source-mips-2.2.14-5S_NF.i386.rpm
   linux-source-ppc-2.2.14-5S_NF.i386.rpm
   linux-source-sparc-2.2.14-5S_NF.i386.rpm
   linux-source-sparc64-2.2.14-5S_NF.i386.rpm

   Install all the packages with the command:

   rpm -Uhv packagename

2.7 Installation with ServeRAID

In this section we describe how to install Caldera OpenLinux products on the xSeries and 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 369.

Stop

Before installing Caldera OpenLinux eServer 2.3 on the xSeries and 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.

Caldera OpenLinux eServer distribution Version 2.3 supports the ServeRAID SCSI controller. The driver version included in the Caldera OpenLinux eServer 2.3 is 3.60.02. This driver release can be only used with the same firmware release. If you are using the ServeRAID 3(L, H or HB) you need to flash the BIOS/firmware of your controller to the 3.50C level which is available on the IBM support site:

`http://www.pc.ibm.com/support`

Note

If your ServeRAID BIOS/firmware is higher than 3.50C you can force downgrade in the BIOS/firmware utility by pressing CTRL+F.

If you are preparing to install on the system with ServeRAID adapter 4(L,M or H) follow the instructions in 2.7.3, “Installing the Caldera OpenLinux on the ServeRAID 4” on page 45, before continuing with installation.

If you are preparing to install on the system with ServeRAID adapter 4(L,M or H) on the system with multiple PCI buses, follow the instructions in 2.7.4, “Installation with multiple PCI buses and ServeRAID 4” on page 48, before continuing with installation.

To install the operating system, follow the procedure in 2.5, “Basic Linux installation” on page 5. Your logical disk drives defined in the RAID array will appear as SCSI drives in the installation program.

Note

Upgrade the drivers and BIOS/firmware of the ServeRAID to the latest level after installation.

After you have installed the system get the latest drivers and utilities for RAID administration from:

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

From that site you can download the following files:

- ips-440.tgz: this file contains the kernel patch for the 2.2.x kernel, which enables the support for IBM ServeRAID adapter in those kernels.
- ipsutils.rpm: this file contains the Linux utilities for IBM ServeRAID SCSI adapter.
- RaidMan.rpm: this file contains the Linux ServeRAID Manager, which can be used to locally or remotely configure and monitor the ServeRAID controller used in Linux installation through the graphical user interface.

This 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 the Linux on the CD are in the \programs\linux directory.

2.7.1 Recompiling the latest kernel with a new ServeRAID driver

In this section we explain how to compile the latest available kernel with the latest available ServeRAID driver. After this you will be able to use the latest kernel with the latest ServeRAID driver.

Because the Caldera OpenLinux eServer distribution 2.3 uses the ServeRAID driver Version 3.60.02 and this driver does not work with the ServeRAID BIOS 4.x, you need to recompile kernel with the support for the latest ServeRAID driver. Also for ServeRAID adapters 4H you must use at least driver Version 4.00 and for 4L/M you must have at least driver Version 4.20.

Note

The ServeRAID BIOS 4.x is only supported with ServeRAID driver 4.x. For ServeRAID adapters 4L/M you must use at least ServeRAID driver 4.20 and for ServeRAID adapter 4H you must use at least ServeRAID driver 4.00.

In our case we will be using the latest available kernel source 2.2.17. This source tree already includes the ServeRAID driver 4.00.06. You can get the source tree from:

```
http://www.gz.us.kernel.org/pub/linux/kernel/v2.2/
```

After downloading the kernel archive linux-2.2.17.tar.bz2 to /tmp directory create the directory for this source tree in /usr/src directory of your Linux installation with the command:

```
mkdir /usr/src/linux-2.2.17
```

Move to the directory and extract the source tree with the commands:

```
cd /usr/src/linux-2.2.17
bzip2 -cd /tmp/linux-2.2.17.tar.bz2 | tar xfv -
```

The source code will be extracted into the /usr/src/linux-2.2.17/linux. Before starting to compile the kernel you must also install the Linux SRPM program from the second CD of your Caldera OpenLinux eServer 2.3 distribution. You can do this by following these commands after inserting the second CD into the CD-ROM:

```
mount /mnt/cdrom
rpm -i /mnt/cdrom/col/sources/SRPMS/linux-2.2.14-1S.src.rpm
```

Copy the original kernel config file to your source tree using the command:

```
cp /usr/src/OpenLinux/SOURCES/linux.defconfig.i586.modular
/usr/src/linux-2.2.17/linux/.config
```

Unpack the driver you downloaded to /tmp directory to the /usr/src/linux-2.2.17 directory with commands:

```
cd /usr/src/linux-2.2.17
tar zxf /tmp/ips-440.tgz
```

Copy the latest drivers to your source tree with the commands:

```
cp ips.c /usr/src/linux-2.2.17/linux/drivers/scsi/ips.c
cp ips.h /usr/src/linux-2.2.17/linux/drivers/scsi/ips.h
```

Move to the source tree directory and recompile the kernel with the commands:

```
cd /usr/src/linux-2.2.17/linux
make menuconfig
```

Note

Do not make any changes in configuration with the menuconfig tool. Simply exit the tool and save the configuration to reflect the latest kernel layout.

```
make dep
make bzImage modules modules_install
```

After compiling the new kernel we copy it into boot directory with the command:

```
cp /usr/src/linux-2.2.17/linux/arch/i386/boot/bzImage
/boot/vmlinuz-2.2.17
```

We also need to copy the System.map file to the /boot directory with the command:

```
cp /usr/src/linux-2.2.17/linux/System.map /boot/System.map-2.2.17
```

Now we need to create the initial RAM disk for our new compiled kernel. First we make a copy of the original initrd file with command:

```
cp /boot/initrd-2.2.14.gz /boot/initrd-2.2.17.gz
```

Then we uncompress the file and mount it to the directory:

```
cd /boot
gzip -d initrd-2.2.17.gz
mkdir /mnt/initrd
mount -o loop initrd-2.2.17 /mnt/initrd
```

Now we go to the /mnt/initrd directory. You will see that it looks like a small file system. The modules are in the linux/modules directory. Move to the directory and execute the following command:

```
cd /mnt/initrd/linux/modules
ls
```

You will see the output similar to Figure 35.

```
[root@nf3500a col]# cd /mnt/initrd/linux/modules/
[root@nf3500a modules]# ls
aic7xxx.o  ips.o  loop.o  scsi_mod.o  sd_mod.o
[root@nf3500a modules]#
```

Figure 35. Modules in initrd

We need to copy all the same modules from our new compiled kernel. For example, copy `ips.o` and `loop.o` with the commands:

```
cp /usr/src/linux-2.2.17/linux/modules/ips.o
cp /usr/src/linux-2.2.17/linux/modules/loop.o
```

After copying all the required modules unmount the initrd file and compress it with the commands:

```
cd /boot
umount /mnt/initrd
gzip -v9 initrd-2.2.17
```

2.7.1.1 Booting with LILO

If you are using LILO for booting you need to tell the LILO (Linux Loader) where to find your new kernel. If you open the LILO configuration file /etc/lilo.conf, you will see your existing configuration, which will be similar to:

```
image = /boot/vmlinuz-pc97-2.2.14-modular
        label  = linux
        root   = /dev/sda3
        vga    = 274
        read-only
        append = "debug=2 noapic nosmp"
        initrd = /boot/initrd-2.2.14.gz
```

Make a copy of it and change it to reflect the new kernel image:

```
image = /boot/vmlinuz-2.2.17
        label  = linux2217
        root   = /dev/sda3
        vga    = 274
        read-only
        append = "debug=2 noapic nosmp"
        initrd = /boot/initrd-2.2.17.gz
```

As you can see we changed the image, label and initrd names. You also need to change the `default` statement to point to the label of new kernel:

```
default = linux2217
```

After saving the file, run the following command to update the boot record:

```
lilo
```

If you have problems with the corrupted text mode after you switch to the text console with Ctrl-Alt-F1-6, you should replace the `vga = 274` option from the kernel configuration with `vga = normal`, and do not forget to rerun LILO after that.

Congratulations! You just created a new working kernel. Restart the server and enjoy the new kernel.

2.7.2 Recompiling the new ServeRAID driver into an existing kernel

In this section we will explain how to create an updated version of ServeRAID driver on the existing kernel source that comes with your Caldera OpenLinux distribution.

Because the Caldera OpenLinux eServer distribution 2.3 uses the ServeRAID driver Version 3.60.02 and this driver does not work with the ServeRAID BIOS 4.x you need to recompile the kernel with the support for the latest ServeRAID driver. Also for the ServeRAID adapters 4H you must use at least driver Version 4.00 and for 4L/M you must use at least driver Version 4.20.

Note

The ServeRAID BIOS 4.x is only supported with the ServeRAID driver 4.x. For ServeRAID adapters 4L/M you must use at least ServeRAID driver 4.20 and for ServeRAID adapter 4H you must use at least ServeRAID driver 4.00.

Before starting to compile the new ServeRAID driver module you must also install the Linux SRPM from the second CD of your Caldera OpenLinux eServer 2.3 distribution. You can do this by following these commands after inserting the second CD into the CD-ROM:

```
mount /mnt/cdrom
rpm -i /mnt/cdrom/col/sources/SRPMS/linux-2.2.14-1S.src.rpm
umount /mnt/cdrom
```

You also need to check if the following packets are installed:

```
linux-source-common,
linux-source-i386.
```

You can do that by executing the following commands:

```
rpm -q linux-source-common
rpm -q linux-source-i386
```

The result should be similar to Figure 36.

```
[root@nf3500a col]# rpm -q linux-source-common
linux-source-common-2.2.14-4
[root@nf3500a col]# rpm -q linux-source-i386
linux-source-i386-2.2.14-4
[root@nf3500a col]#
```

Figure 36. Checking the source code install

If the code is not installed, then insert the first CD of your Caldera OpenLinux distribution and install the packages with these commands:

```
mount /mnt/cdrom
rpm -i /mnt/cdrom/col/install/RPMS/linux-source-common-2.2.
14-1S.i386.rpm
rpm -i /mnt/cdrom/col/install/RPMS/linux-source-i386-2.2.14
-1S.i386.rpm
umount /mnt/cdrom
```

Copy the original kernel config file to your source tree using the command:

```
cp /usr/src/OpenLinux/SOURCES/linux.defconfig.i586.modular
/usr/src/linux/.config
```

Unpack the driver you downloaded to the /tmp directory to the /usr/src directory with the commands:

```
cd /usr/src/
tar zxf /tmp/ips-440.tgz
```

Copy the latest drivers to your source tree with the commands:

```
cp ips.c /usr/src/linux/drivers/scsi/ips.c
cp ips.h /usr/src/linux/drivers/scsi/ips.h
```

Move to the source tree directory and recompile the kernel with the commands:

```
cd /usr/src/linux
make menuconfig
```

Note

Do not make any changes in configuration with the menuconfig tool. Simply exit the tool and save the configuration to reflect the latest kernel layout.

```
make dep
make modules
```

After compiling we need to copy the new ServeRAID driver to the modules directory of the running kernel with the command:

```
cp /usr/src/linux/modules/ips.o /lib/modules/2.2.14/scsi/ips.o
```

After copying the new module you need to run the following command to update the modules dependency:

```
depmod -a
```

Now we need to update the initial RAM disk with our new compiled ServeRAID module. First we uncompress the file and mount it to the directory:

```
cd /boot
gzip -d initrd-2.2.14.gz
mkdir /mnt/initrd
mount -o loop initrd-2.2.14 /mnt/initrd
```

Now we go to the /mnt/initrd directory. You will see that it looks like a small file system. The modules are in the linux/modules directory. Move to the directory and execute the following command:

```
cd /mnt/initrd/linux/modules
ls
```

You will see the output similar to Figure 37.

```
[root@nf3500a col]# cd /mnt/initrd/linux/modules/
[root@nf3500a modules]# ls
aic7xxx.o  ips.o  loop.o  scsi_mod.o  sd_mod.o
[root@nf3500a modules]#
```

Figure 37. Modules in initrd

We need to copy the new ServeRAID module with the command:

```
cp /lib/modules/2.2.14/scsi/ips.o .
```

After copying all the required module, unmount the initrd file and compress it with the commands:

```
cd /boot
umount /mnt/initrd
gzip -v9 initrd-2.2.14
```

If you are using LILO(Linux Loader) for booting, then you need to tell LILO that you changed your initrd image. To reflect the changes, simply run LILO with the command:

```
lilo
```

Congratulations! You just updated your kernel with the latest ServeRAID driver. Restart the server and enjoy using ServeRAID.

2.7.3 Installing the Caldera OpenLinux on the ServeRAID 4

To successfully install the Caldera OpenLinux eServer 2.3 on the IBM @server systems with ServeRAID adapter 4 installed you need to create a customized installation disk. For this you will need to install Caldera OpenLinux on some other supported computer. After you installed the system follow these steps to create a customized installation disk:

1. Following the instructions in 2.7.2, “Recompiling the new ServeRAID driver into an existing kernel” on page 42, compile the latest ServeRAID driver 4.40.03 as a module.
2. From the first installation CD-ROM of the Caldera OpenLinux distribution, build the original installation disk executing the commands:
   - Insert the installation CD-ROM:

     ```
     mount /mnt/cdrom
     ```

   - Insert the blank floppy disk:

     ```
     dd if=/mnt/cdrom/col/launch/floppy/install.144 of=/dev/fd0
     ```

3. Mount the floppy and copy the initrd.gz from the floppy disk to /tmp directory with the commands:

   ```
   mount /mnt/floppy
   cp /mnt/floppy/initrd.gz /tmp
   ```

4. Unzip the initrd and mount it as a loop device with the commands:

   ```
   cd /tmp
   gzip -d initrd.gz
   mkdir /mnt/initrd
   ```

```
mount -o loop initrd /mnt/initrd
```

5. Delete some of the SCSI drivers to make room for the new ServeRAID driver and then copy the latest ips.o ServeRAID driver you compiled in step 1 into the initrd with the commands:

```
rm /mnt/initrd/lib/modules/2.2.14/scsi/advansys.o
rm /mnt/initrd/lib/modules/2.2.14/scsi/fdomain.o
rm /mnt/initrd/lib/modules/2.2.14/scsi/ncr53c8xx.o
cp /usr/src/linux/modules/ips.o /mnt/initrd/lib/modules/2.2.14/scsi
```

6. Now you need to edit the /etc/pcidrivers file in initrd to include the ID of the ServeRAID 4 adapters. This is necessary if you want that installation program recognizes your ServeRAID 4 adapter. To do this open the file /mnt/initrd/etc/pcidrivers in your favorite editor for example:

```
vi /mnt/initrd/etc/pcidrivers
```

   Find the following section:

```
driver ips
1014002e
```

   After the line with `1014002e` add the line with `101401bd`.. After the change this section should look like this:

```
driver ips
1014002e
101401bd
```

7. Unmount the initrd file, compress it and copy it back to the installation floppy with the commands:

```
cd /tmp
umount /mnt/initrd
gzip -v9 initrd
cp initrd.gz /mnt/floppy
umount /mnt/floppy
```

After creating the customized installation disk, insert the first installation CD into the server, and boot your IBM @server with customized installation disk. The installation program will now detect the ServeRAID 4 adapter. Continue installation following the instructions in 2.5, “Basic Linux installation” on page 5. When the installation finishes and the system is started switch to the text console with the CTRL+ALT+F1 key combination and log on as a root. Now we need to update the initrd file of the new installation. You will do this following this steps:

1. Mount the customized installation disk you used for installation and copy the initrd from it to the /tmp directory with the commands:

   ```
   mount /mnt/floppy
   cp /mnt/floppy/initrd.gz /tmp
   ```

2. Uncompress the initrd you just copied and mount it as a loop device with commands:

   ```
   cd /tmp
   gzip -d initrd.gz
   mkdir /mnt/initrd
   mount -o loop initrd /mnt/initrd
   ```

3. Uncompress the installed initrd from /boot directory and mount it as a loop device with commands:

   ```
   cd /boot
   gzip -d initrd-2.2.14.gz
   mkdir /mnt/newinitrd
   mount -o loop initrd /mnt/newinitrd
   ```

4. Copy the ips.o driver from the installation disk to the installed initrd and to modules directory with the commands:

   ```
   cp /mnt/initrd/lib/modules/2.2.14/scsi/ips.o
   /mnt/newinitrd/linux/modules
   cp /mnt/initrd/lib/modules/2.2.14/scsi/ips.o /lib/modules/2.2.14/scsi
   ```

5. Unmount both initrd and compress the installed initrd with the commands:

   ```
   cd /boot
   umount /mnt/initrd
   umount /mnt/newinitrd
   gzip -v9 initrd-2.2.14
   ```

6. Run LILO to reflect the changes in the initrd file with the command:

   ```
   lilo -v
   ```

7. Reboot the system and enjoy the performance of the ServeRAID 4 adapter.

Congratulations! You have installed Caldera OpenLinux eServer 2.3 on the IBM @server with ServeRAID 4 installed.

You can check if the correct driver is loaded by executing the command:

```
cat /proc/scsi/ips/2
```

The last number depends on your configuration and it could be 1 or something else. If the correct driver is loaded you should see output similar to Figure 38.

```
[root@x230 col]# cat /proc/scsi/ips/2

IBM ServeRAID General Information:

        Controller Type                 : ServeRAID 4M
        Memory region                   : 0xf7ffc000 (8192 bytes)
        Shared memory address           : 0x8803e000
        IRQ number                      : 20
        BIOS Version                    : 4.40.03
        Firmware Version                : 4.40.03
        Boot Block Version              : 4.40.03
        Driver Version                  : 4.40.03
        Max Physical Devices            : 30
        Max Active Commands             : 128
        Current Queued Commands         : 0
        Current Active Commands         : 0
        Current Queued PT Commands      : 0
        Current Active PT Commands      : 0
```

Figure 38. ServeRAID 4 status report

2.7.4 Installation with multiple PCI buses and ServeRAID 4

To successfully install the Caldera OpenLinux eServer 2.3 on the IBM multiple PCI buses and ServeRAID adapter 4, with the firmware/BIOS level above 4.00.06 installed, you need to create a customized installation disk. For this you will need to install the Caldera OpenLinux using an updated CD for the IBM server on some other supported computer. After you have installed the system, follow these steps to create a customized installation disk:

1. Apply the latest packages for the IBM systems following the instructions in 2.6, “Installation on the systems with multiple PCI buses” on page 35 in step 3.

Note

You have to use an updated installation and source CD for the following steps. Instead of using the linux-2.2.14-3S.src.rpm from an updated source CD you should use the file linux-2.2.14-5S_NF.src.rpm, from the Caldera Systems FTP site:

`ftp://ftp.calderasystems.com/pub/eServer/updates/2.3/extras/IBM/SRPMS`

2. Follow the instructions in 2.7.2, “Recompiling the new ServeRAID driver into an existing kernel” on page 42, but before executing the command `make dep`, follow these instructions:

   a. Execute the command `make menuconfig`. You will see a screen similar to Figure 39.

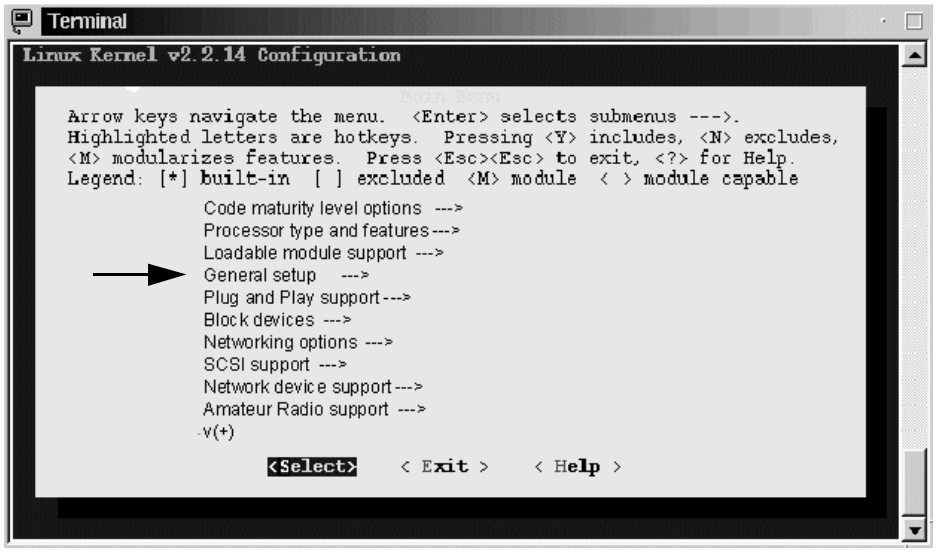

Figure 39. Make menuconfig window

   b. Select **General Setup** and you will see a screen similar to Figure 40.

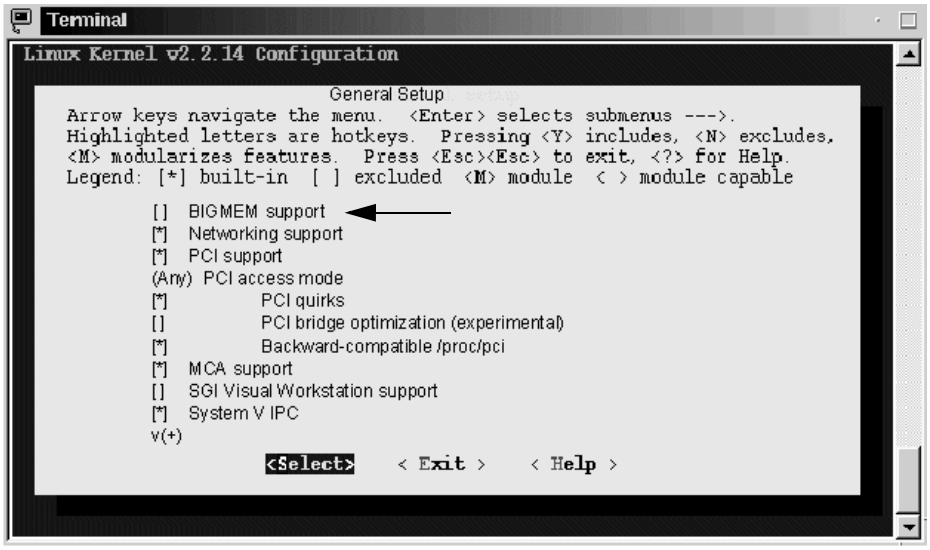

Figure 40. General setup window

c. Select the option **BIGMEM support,** as shown in Figure 40.

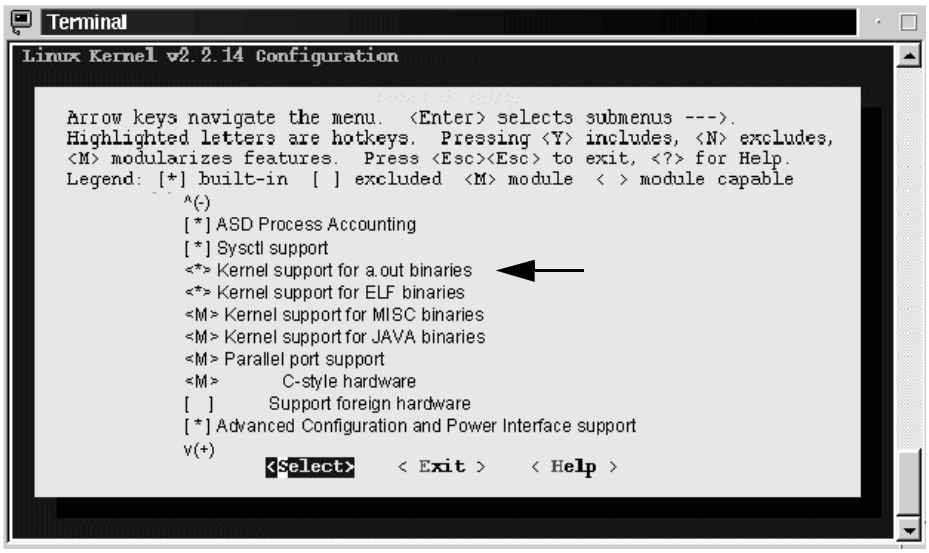

Figure 41. a.out support

d. Select the option **Kernel support for a.out binaries**.Figure 41.

e. Save the configuration file and exit the configuration tool.

f. Because the configuration file from an updated installation and source CD already includes the ServeRAID and SCSI controllers compiled into kernel, you need to run the following commands now:

```
make clean
make dep
make bzImage
```

3. Obtain the special installation disk for IBM systems with multiple PCI buses from CalderaSystems FTP site:

```
ftp://ftp.calderasystems.com/pub/eServer/updates/2.3/extras/IBM/launch/
floppy/install-2.2.14-244-syslinux.img
```

4. Build the floppy disk with the command:

```
dd if=install-2.2.14-244-syslinux.img of=/dev/fd0
```

5. Mount the floppy and copy the new kernel you just compiled on the floppy disk with the commands:

```
mount /mnt/floppy
cp /usr/src/linux/arch/i386/boot/bzImage /mnt/floppy/vmlinuz
umount /mnt/floppy
```

After creating the customized installation disk, insert the first updated installation CD into the server, and boot your IBM @server with a customized installation disk. The installation program will now detect the ServeRAID 4 adapter. Continue installation following the instructions in 2.5, “Basic Linux installation” on page 5. When the installation finishes and the system is started, switch to the text console by pressing the CTRL+ALT+F1 key combination and log on as a root**.** Now we need to copy our kernel to the new installation:

1. Mount the customized installation disk you used for the installation and copy the kernel to the /boot directory with the commands:

```
mount /mnt/floppy
```

```
cp /mnt/floppy/vmlinuz /boot/vmlinuz-pc97-2.2.14-modular
```

2. Run the LILO to reflect the change of the new kernel with the command:

```
lilo -v
```

Congratulations! You have installed Caldera OpenLinux eServer 2.3 on the IBM @server with multiple PCI buses and ServeRAID 4 installed.

After installation apply the latest packages for the IBM systems following the instructions in 2.6, “Installation on the systems with multiple PCI buses” on page 35 in step 3.

2.7.5 Installing ipsutils.rpm

Before installing the ipsutils package you need to get the latest versions of the glibc packages from the Caldera Systems FTP site:

```
ftp://ftp.calderasystems.com/pub/updates/eServer/2.3/current/RPMS
```

The version of the packages has to be at least:

```
glibc-devel-static-2.1.3-4S.i386.rpm
glibc-devel-2.1.3-4S.i386.rpm
glibc-2.1.3-4S.i386.rpm
glibc-localedata-2.1.3-4S.i386.rpm
```

After you download the required packages, install them with the commands:

```
rpm -Uhv --nodeps glibc-2.1.3-4S.i386.rpm
rpm -Uhv --nodeps glibc-devel-2.1.3-4S.i386.rpm
rpm -Uhv glibc-devel-static-2.1.3-4S.i386.rpm
rpm -Uhv glibc-localedata-2.1.3-4S.i386.rpm
```

To successfully 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. Caldera OpenLinux uses RPM for installing packages, so the RPM utility is already installed on your system. To check if your RPM utility is working, open a terminal window in the KDE graphic environment and execute this command:

```
rpm --version
```

Your RPM is working if your output looks something like this:

```
RPM version 2.5.5.OL
```

More GUI-oriented users will probably want to use the kpackage tool, which is part of the KDE graphic environment and is used for installing packages. But if you try to install the ipsutils.rpm with this tool, you will see a window similar to Figure 42.

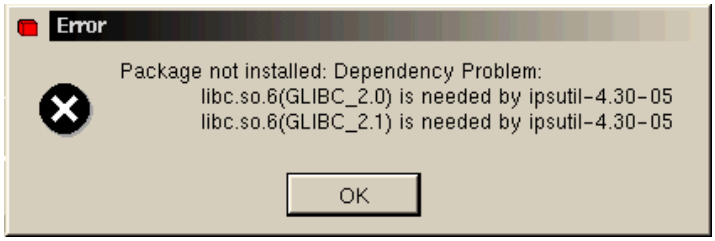

Figure 42. Error installing ipsutils.rpm with kpackage

Experienced users will recognize that this is a script, but even if you try to install without enabling the **Check Dependencies** option, as you can see in Figure 43, you will still get the same error.

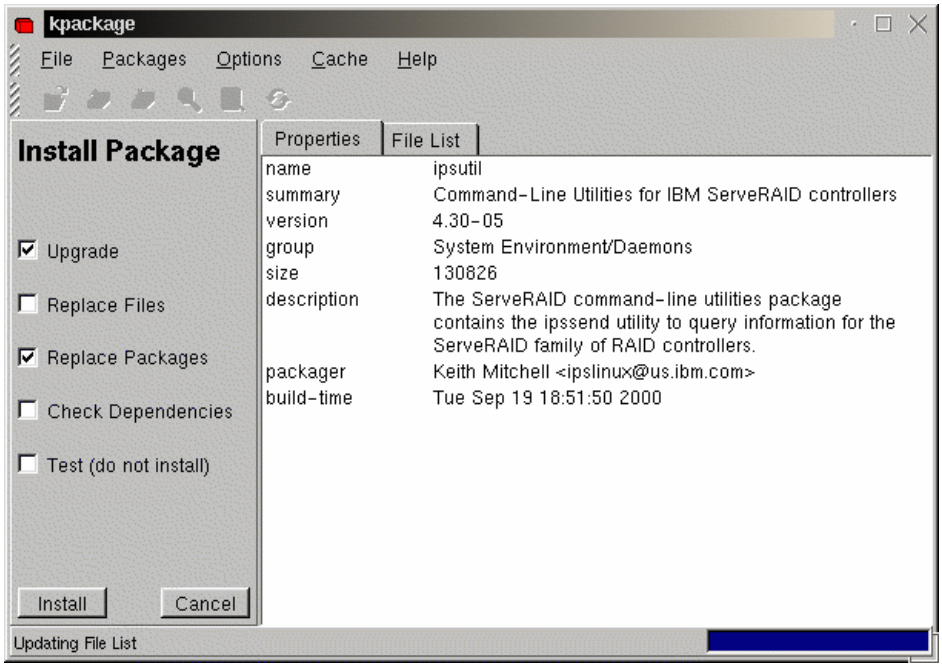

Figure 43. Installing without dependencies in kpackage

This is because ipsutils.rpm displays the copyright message before installing. So the only way to install the package is to execute the following command from a terminal:

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

Note

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

```
mount /mnt/cdrom

cp /mnt/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.prm 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 will see output similar to Figure 44.

```
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 44. 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 ServeRAID controller efficiently.

2.7.6 The ipssend commands

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

2.7.6.1 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 45.

```
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 ST39175LO4303AL0A27C
```

Figure 45. 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:

```
ispsend getconfig 1 ad
```

You will see output similar to Figure 46.

```
[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 46. 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 47.

```
[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 47. 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 48.

```
[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 ST39175LO4303AL0A27C
      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 ST39175LO4303AL09YSS
      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 ST39175LO4303AL0A2QK
      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 48. Executing ipssend getconfig 1 pd

2.7.6.2 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>
```

The parameters are explained in Table 2.

Table 2. getstatus command parameters

Parameter	Description
Controller	Number of controller (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 49.

```
[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 49. Executing ipssend getstatus 1

If the ServeRAID controller is in the middle of rebuilding a drive, you will see output similar to Figure 50.

```
[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 50. Executing ipssend getstatus 1 during rebuilding of a drive

2.7.6.3 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 3.

Table 3. devinfo command parameters

Parameter	Description
Controller	Number of controller (1 to 12)
Channel	Channel of device (1 to 3)

Parameter	Description
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 51.

```
[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 51. 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 52.

```
[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 52. Executing ipssend devinfo 1 1 2 during rebuilding of a drive

2.7.6.4 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 4.

Table 4. 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 53.

```
[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 53. 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 54.

```
[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 54. Executing ipssend hsrebuild 1 on

2.7.6.5 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 5.

Table 5. 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! For example, 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 the 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 55.

```
[root@nf350a /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 55. Executing ipssend setstate 1 1 0 rdy

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

```
ipssend getconfig 1 pd
```

2.7.6.6 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 6.

Table 6. 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.

2.7.6.7 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 7.

Table 7. 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 56.

```
[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 56. 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 57.

```
[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 57. Executing ipssend unattended 1 on

2.7.6.8 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 8.

Table 8. 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 ServeRAID 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 58.

```
[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 58. Executing ipssend rebuild 1 1 2 1 0

2.7.7 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 50 on page 60.

If the rebuild is not completed successfully, you will see output similar to Figure 59.

```
[root@nf350Oa /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 59. Failed rebuild

2.7.8 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@nf350Oa /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 60. Forced rebuild of the defunct drive

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

```
ipssend getstatus 1
```

You will see output similar to Figure 50 on page 60.

2.7.9 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've realized 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 originally 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 61.

```
[root@nf350Ua /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 ST39175LO4303AL0A27C
      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 ST39175LO4303AL09YSS
      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 ST39175LO4303AL0A2QK
      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 61. 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 62.

```
[root@nf350Oa /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 62. 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 63.

```
[root@nf350Oa /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 (SBY)
         Size (in MB)/(in Sectors): 0/0
         Device ID                : IBM     EXP200  10D792063452
Command Completed Successfully.
```

Figure 63. 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 in 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 64.

```
[root@nf350a /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 64. 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 65.

```
[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 65. 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 66.

```
[root@nf350Oa /root]# ipssend setstate 1 1 3 hsp

Found 1 IBM ServeRAID Controller(s).
Set Device State has been initiated for Controller 1...
Command Completed Successfully.
```

Figure 66. 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 67.

```
[root@nf350Oa /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 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                    : Hot Spare (HSP)
         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 67. 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.

2.7.10 Using the ServeRAID Manager utility

The ServeRAID Manager for Linux helps 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 with 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, you install it with the command:

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

After invoking this command you will see the screen similar to Figure 68.

```
[root@x220 col]# rpm -ihv RaidMan-4.40-03.i386.rpm
failed dependencies:
        libstdc++-libc6.1-1.so.2 is needed by RaidMan-4.40-03
        libc.so.6(GLIBC_2.0) is needed by RaidMan-4.40-03
        libc.so.6(GLIBC_2.1) is needed by RaidMan-4.40-03
        libm.so.6(GLIBC_2.1) is needed by RaidMan-4.40-03
[root@x220 col]#
```

Figure 68. Failed install of ServeRAID Manager

To successfully install this version of ServeRAID Manager you need to get the latest versions of the glibc packages from the Caldera Systems FTP site:

```
ftp://ftp.calderasystems.com/pub/updates/eServer/2.3/current/RPMS
```

The version of the packages has to be at least:

```
glibc-devel-static-2.1.3-4S.i386.rpm
glibc-devel-2.1.3-4S.i386.rpm
```

```
glibc-2.1.3-4S.i386.rpm
glibc-localedata-2.1.3-4S.i386.rpm
```

After you download the required packages, install them with the commands:

```
rpm -Uhv --nodeps glibc-2.1.3-4S.i386.rpm
rpm -Uhv --nodeps glibc-devel-2.1.3-4S.i386.rpm
rpm -Uhv glibc-devel-static-2.1.3-4S.i386.rpm
rpm -Uhv glibc-localedata-2.1.3-4S.i386.rpm
```

Now install ServeRAID Manager with the command:

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

Note

Before using this version of the ServeRAID Manager software you must have BIOS/firmware and the driver for the controller 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 into the /etc/inittab file:

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

This will start the agent in every run level. The installation program will also start 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 following directory:

```
/usr/RaidMan
```

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

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

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

Note

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

During the program startup you will see a window similar to Figure 69.

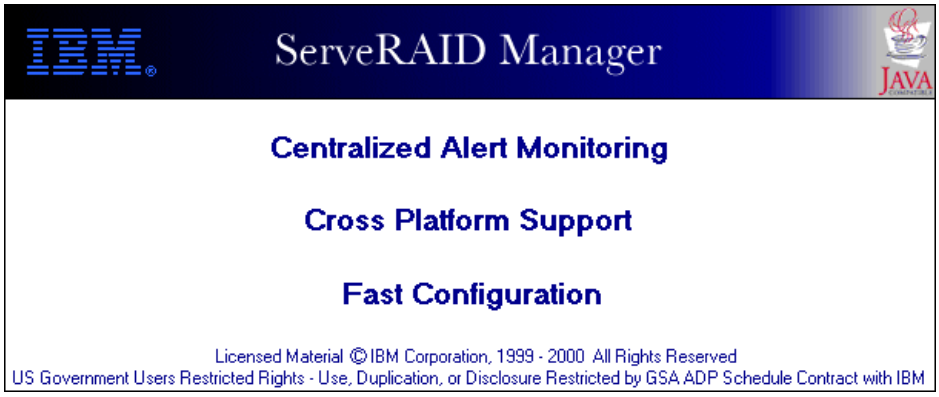

Figure 69. ServeRAID Manager startup

After the program is started you will see a window similar to Figure 70.

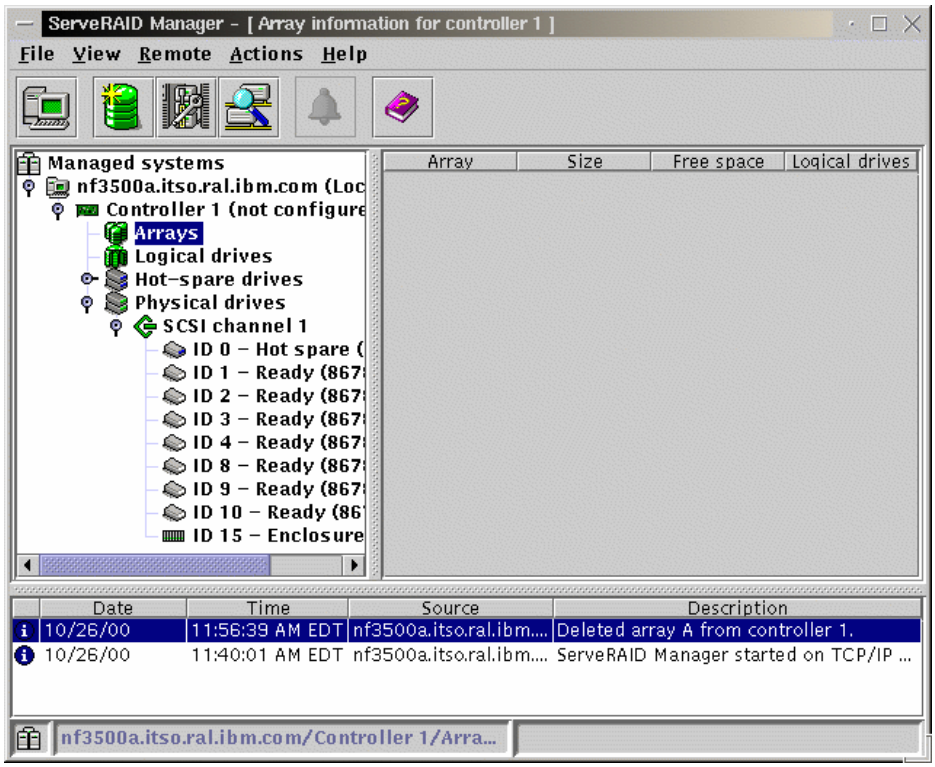

Figure 70. ServeRAID Manager

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

- Menus - in the menus you can access all the functions available
- Icons - icons offers you shortcuts to the most often used functions
- Tree window - here you can see all the systems with 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 about how to use the ServeRAID Manager functions can be found in the online help.

2.7.11 Remote management of ServeRAID adapter

Your server witha ServeRAID controller installed can also be managed remotely. For this you need the following:

1. Install ServeRAID Manager on the server with an installed ServeRAID adapter as we described in 2.7.10, “Using the ServeRAID Manager utility” on page 75. Do not forget to enable the ServeRAID Manager agent running as a 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 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 properly enable the access to your server. By default the security is enabled after the installation of the ServeRAID Manager.

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

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

2. From the Actions menu select **Configure ServeRAID agent->Security** as you can see in Figure 71.

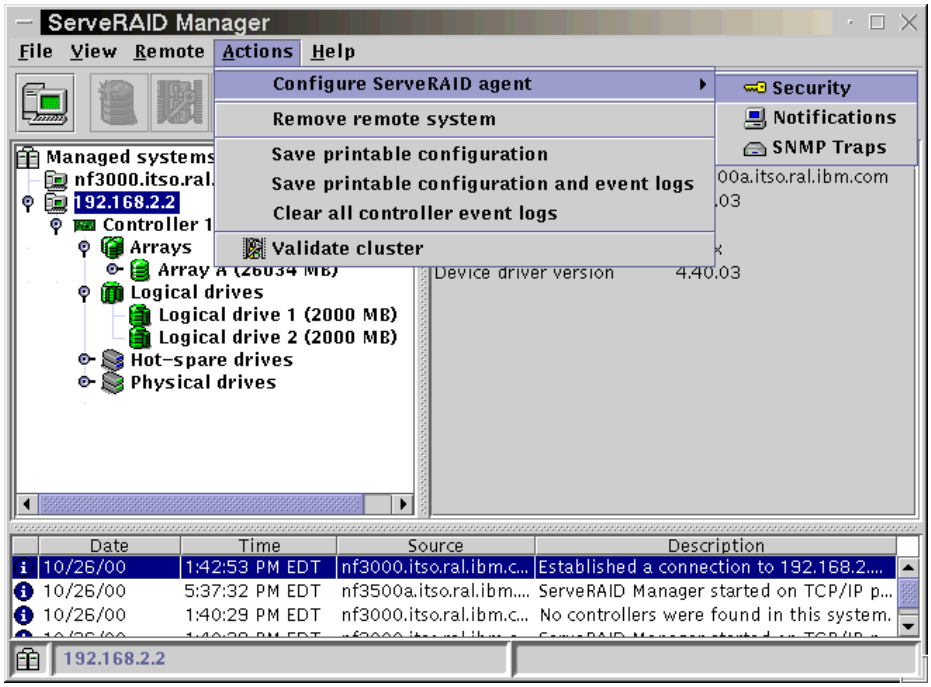

Figure 71. Starting security configuration

You will see a window similar to Figure 72.

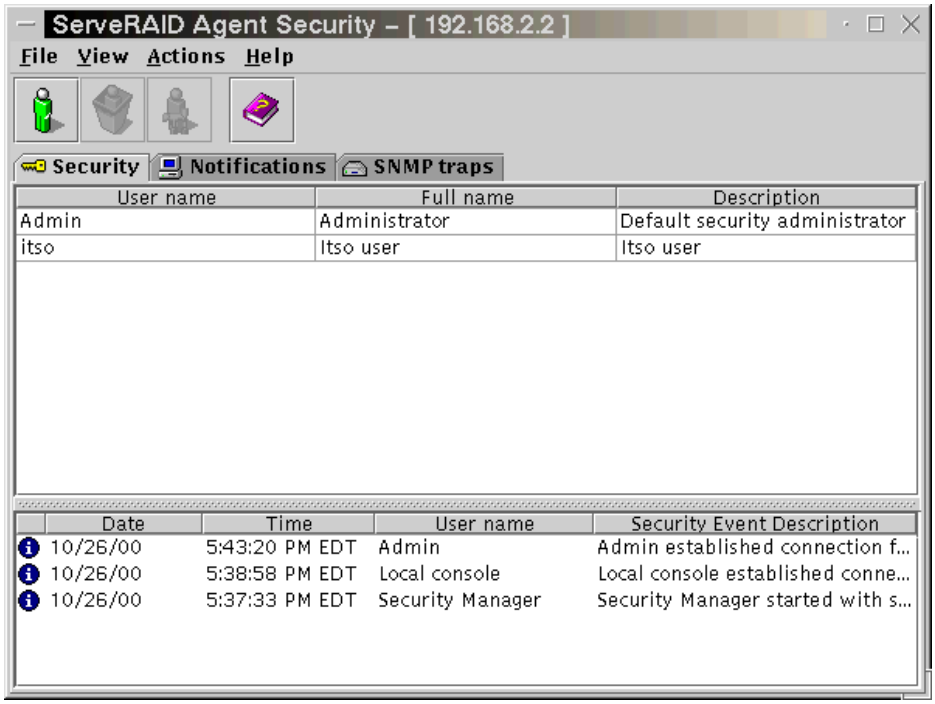

Figure 72. Security window

3. Double-click **Admin** (this is a built-in user that cannot be removed) and you will see the window similar to Figure 73.

Figure 73. Changing the user

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

6. Install the ServeRAID Manager on the server that will 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. With the command:

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

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

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

8. By using this command from the command prompt, you restart the ServeRAID agent:

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

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

You can connect to a remote server from a management workstation by selecting **Remote -> Add remote system** as you can see in Figure 74.

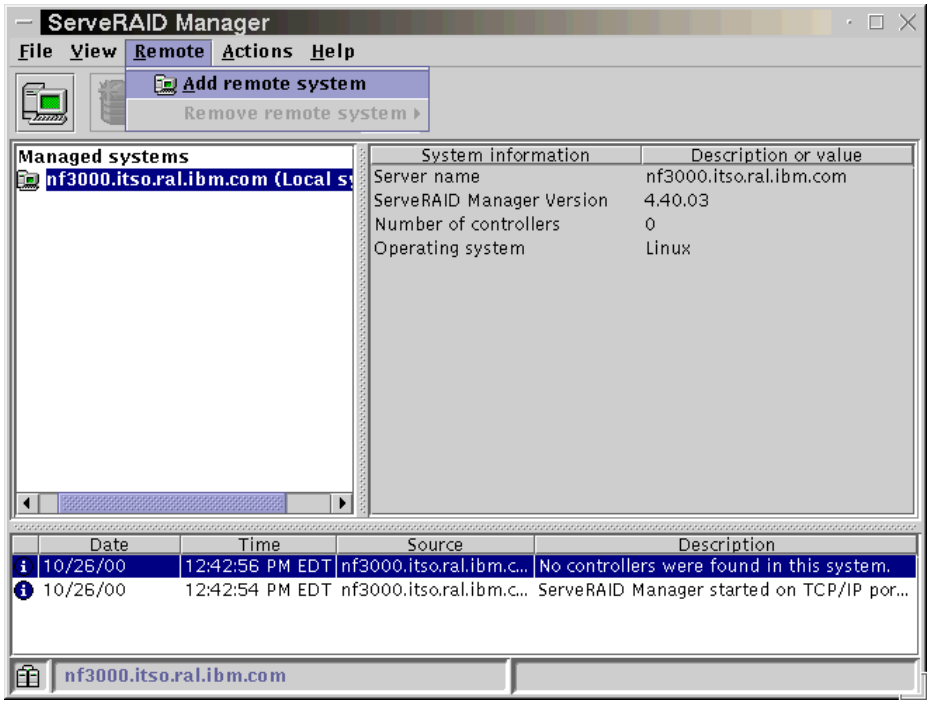

Figure 74. Accessing the remote server

You will see a window similar to Figure 75.

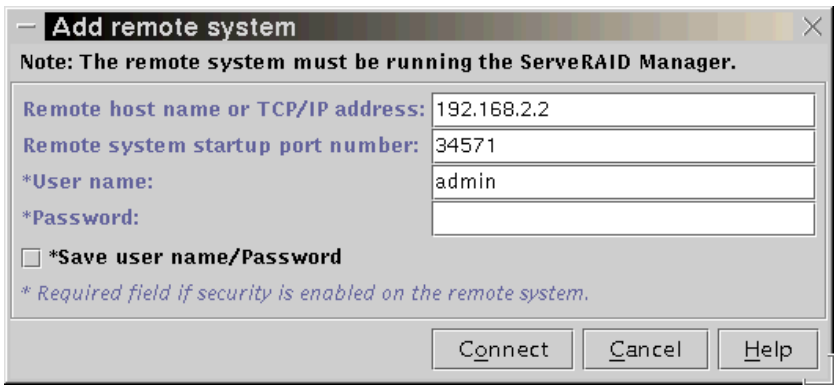

Figure 75. Remote system

Type in the necessary data and click **Connect.** After the system is connected you will a window similar to Figure 76.

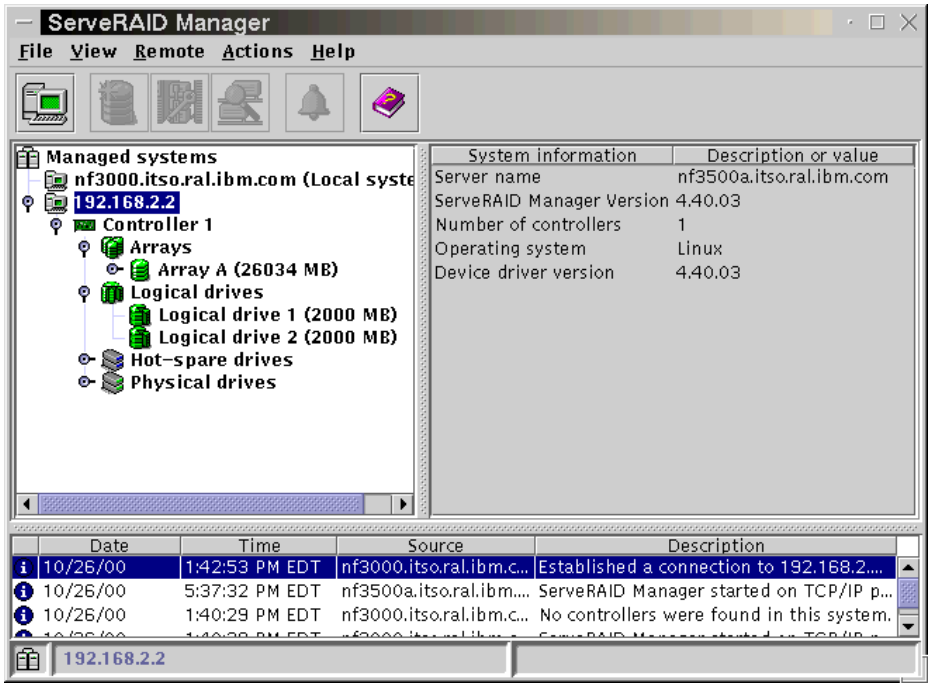

Figure 76. Remote system after connection

Now you can start managing the remote ServeRAID adapter.

2.8 Installing and configuring token-ring network cards

In Caldera OpenLinux eServer 2.3 there is no way to graphically add and configure token-ring cards. To successfully install and configure token-ring cards, follow these steps:

1. Install a token-ring card (PCI or ISA) in a free slot in your IBM Netfinity server.
2. Start the system and from COAS start Kernel tools as described in 3.20, “Manipulating kernel modules” on page 127. To skip the probing of new hardware, click **Done**. When the tool is started, you will see a window similar to Figure 77.

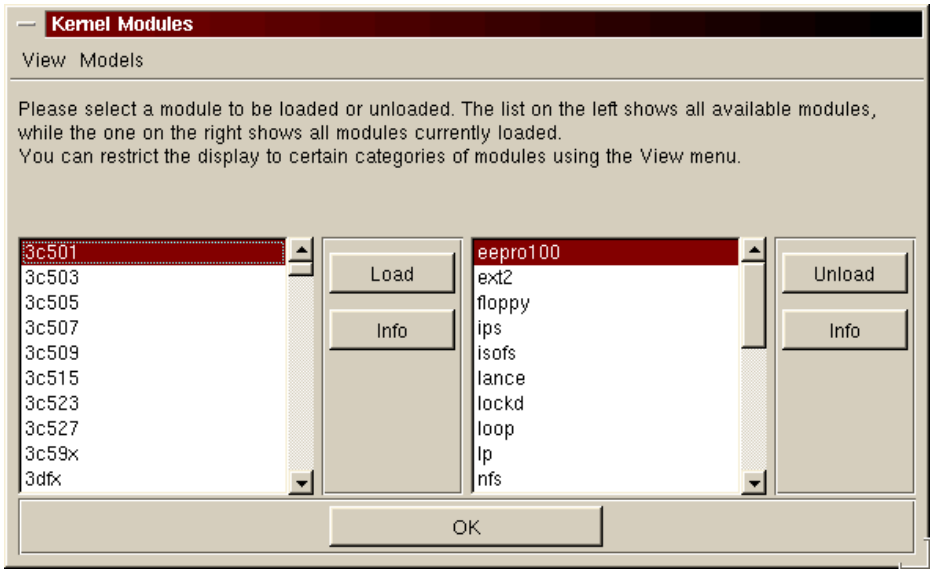

Figure 77. Loading the module for token-ring card

3. Load the appropriate module for your card:
   - **olympic** - this is the module for a PCI token-ring network card.
   - **lanstreamer** - this is the module for a IBM PCI LANstreamer card.
   - **ibmtr** - this is the module for an ISA token-ring network card.

   In our example we used the olympic module for the PCI token-ring network card. Select the module and click **Load,** and you will see a window similar to Figure 78.

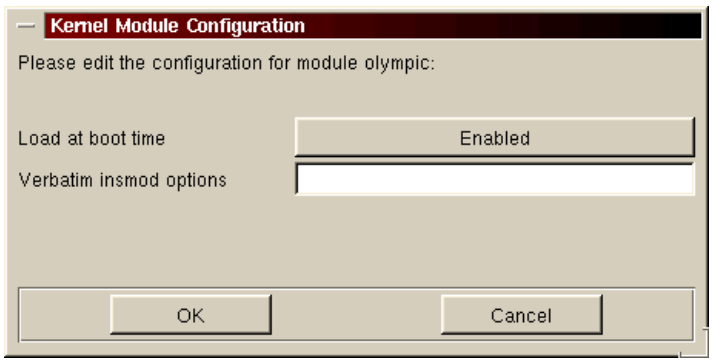

Figure 78. Loading the module for PCI token-ring network card

Note

When using the IBM Turbo 16/4 ISA adapter you must set it to auto 16/4 compatibility mode. This is done by using the LANAID diskette and running this command:

```
LANAIDC /fast=auto16
```

If you are not sure, run `LANAIDC /view`.

4. Click **OK** to load the module. If the module is loaded successfully the window will close and you will see a green light on the back of the network card.
5. You need to create a network script for the token-ring network card. In our example we are assuming that this is the first token-ring network card in your server. With your favorite editor, create the /etc/sysconfig/network-scripts/ifcfg-tr0 file. It should look like this:

```
DEVICE=tr0
IPADDR=9.24.104.202
REMIP=0.0.0.0
NETMASK=255.255.255.0
BROADCAST=9.24.104.255
MTU=1500
GATEWAY=9.24.104.1
ONBOOT=yes
```

   You need to adapt all addresses to your configuration.

Note

Caldera OpenLinux does not support DHCP for the token-ring cards. If you want to have DHCP running on your token-ring card you need to download the DHCP client that supports the token-ring cards from `http://www.isc.org/` and configure the /etc/sysconfig/network-scripts/ifup-dhcp script to use the new client.

   If you want to use DHCP with your network card then you have to add an additional line to the /etc/sysconfig/network-scripts/ifcfg-tr0 file:

```
DYNAMIC=dhcp
```

   You can see the example of a modified /etc/sysconfig/network-scripts/ifup-chcp script in Appendix F, “Modified ifup-dhcp file” on page 399.

In case you have more than one token-ring card in your server, you need to create the configuration file for each of them. The file names will be `ifcfg-tr0`, `ifcfg-tr1`, `ifcfg-tr2` and so on.

6. Restart the network with the following commands:

```
/etc/rc.d/init.d/network stop
/etc/rc.d/init.d/network start
```

Congratulations! Your token-ring network card is now working and is ready for some heavy traffic.

Chapter 3. Basic system administration

This chapter will give you an overview of how to perform the most common administrative tasks on a Caldera OpenLinux eServer 2.3 operating system. Most of these tasks can be done with the Caldera Open Administration System (COAS), Caldera's OpenLinux graphical-oriented configuration and administration tool. However, you may still perform these tasks using the command-line tools.

Stop

Be careful when you are editing configuration files on your own. If you edit configuration files with an editor, make sure to maintain the format of the file. If you change the format of a configuration file, COAS may not be able to understand the configuration information and you cannot use COAS for future configuration.

3.1 Log in to the system

Before you can use any Linux system you need to log in to the system. Whenever you start Caldera OpenLinux, you will see a login window similar to Figure 79.

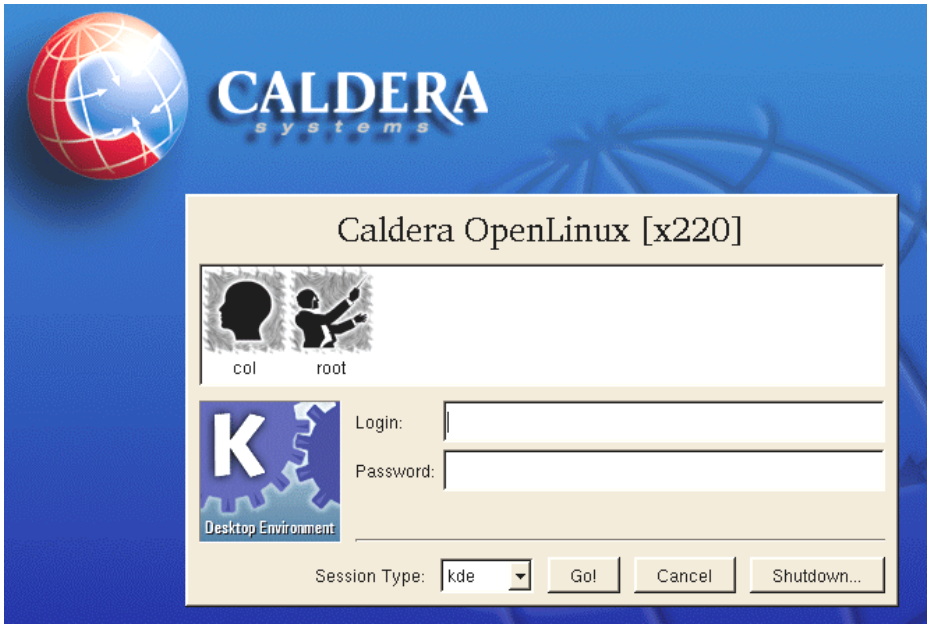

Figure 79. Login window

If you wish to use a text-based user interface, you can press Ctrl+Alt+Fx, where x is the number from 1 to 6, to switch to a text console. For example to switch to console 1, you need to press Ctrl+Alt+F1, and you will see a window similar to Figure 80.

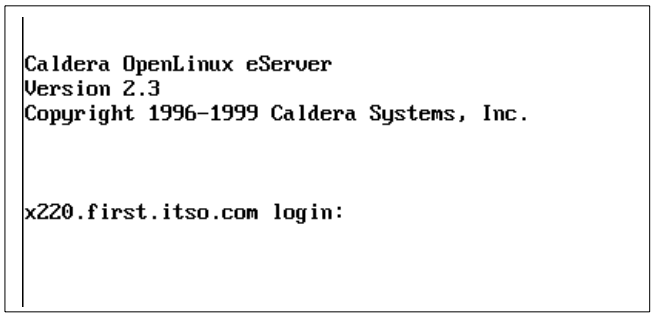

Figure 80. Text-based user interface

If you want to switch back to the graphical interface press Alt-F8**.** This means that you are switching to the console number 8. Caldera OpenLinux uses this console for the graphical user interface. To start working with Caldera OpenLinux, you need to log on with either a graphical or a text-based user

interface. To start the graphical user interface, type in the user name and password in the window shown in Figure 79, and click **Go!** You will see a window similar to Figure 81.

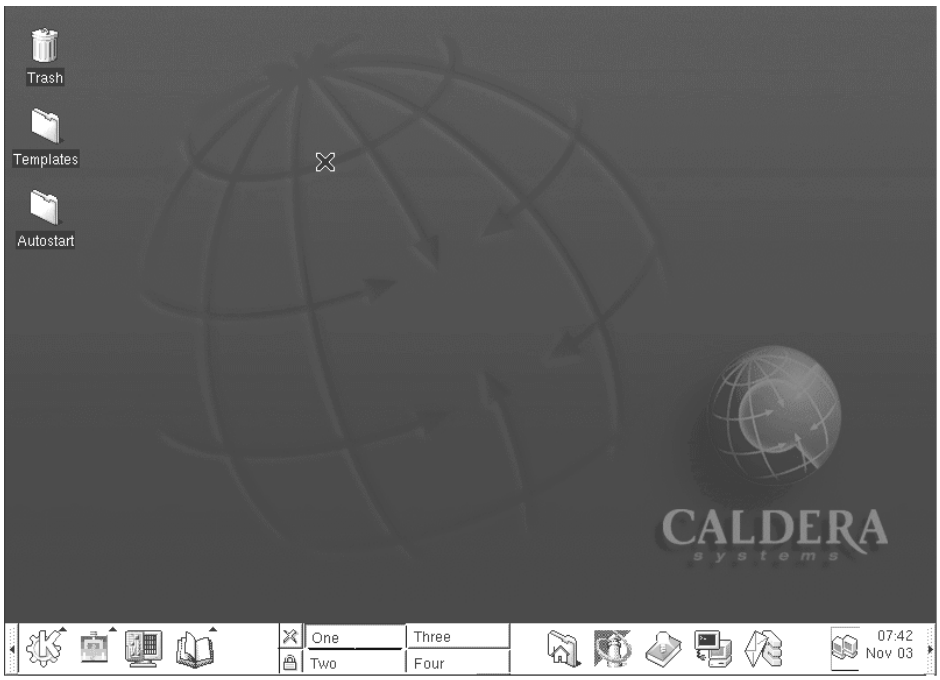

Figure 81. KDE Window Manager

3.2 Using the Window Manager

Once you are logged into the system through the graphical user interface you will see a window similar to Figure 81, which is controlled by the Window Manager. Caldera OpenLinux uses the KDE Window Manager. You can get more information about KDE on:

```
http://www.kde.org
```

At the bottom of the window you can see the toolbar that is used for accessing all available functions. It has pull-down menus, icons and buttons. You can use them for accessing the features of the operating system and applications.

Note

We recommend that you use more than 8bpp color definitions for your XFree86 server setup; otherwise, you will have problems with missing colors when you open more programs.

In the following sections we will describe how to use some basic tools in the graphical environment and especially how to customize your Caldera OpenLinux system by using COAS.

3.3 Getting the X-Windows terminal window

In order to run commands from the command line when you have the GUI Windows-based window in front of you, you need to create a terminal window. You can do this by clicking the icon representing the terminal window, circled in Figure 82.

Figure 82. Starting terminal window

After the terminal window is started, you will see a window similar to Figure 83.

Figure 83. Terminal window in KDE

In this terminal window you can access the system from a command-line prompt as in a text-based interface. The command line prompt gives you more flexibility than menus, but you can do most of the basic things from the menu system. It is a matter of personal choice.

3.4 Accessing COAS - Caldera Open Administration System

All the administration tasks in Caldera OpenLinux are performed through the use of COAS. You can access the COAS tools by clicking the **COAS** icon on the KDE toolbar, circled in Figure 84.

Figure 84. Accessing the COAS tools

After you click the **COAS** icon, you will see a window similar to Figure 85.

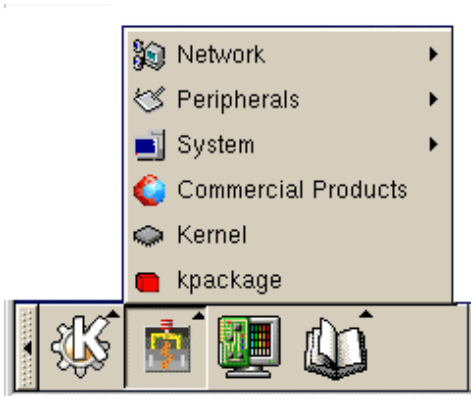

Figure 85. COAS tools

You can see you have several tools available. We will discuss them in the following sections.

3.5 Adding and removing software packages using kpackage

If you want to add or remove software once Caldera OpenLinux is installed or just check if the software is installed, you can do this by using the kpackage tool. You can start kpackage by selecting **kpackage** from the COAS tools menu, as you can see in Figure 86.

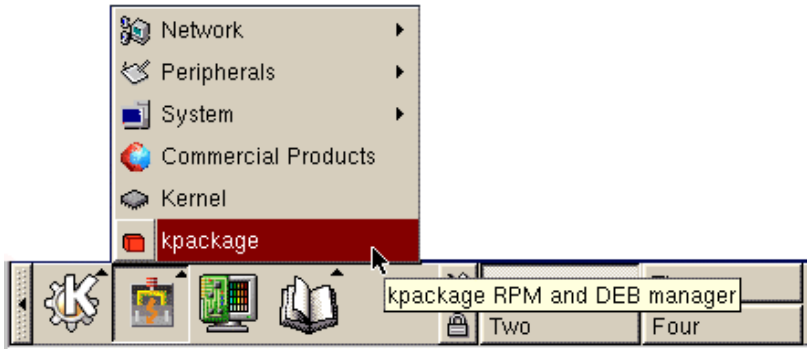

Figure 86. Starting kpackage

When kpackage is started, you will see a window similar to Figure 87.

Figure 87. kpackage

3.5.1 Uninstalling a package

If you want to uninstall a package, select the desired package and click **Uninstall.** You will see a window similar to Figure 88.

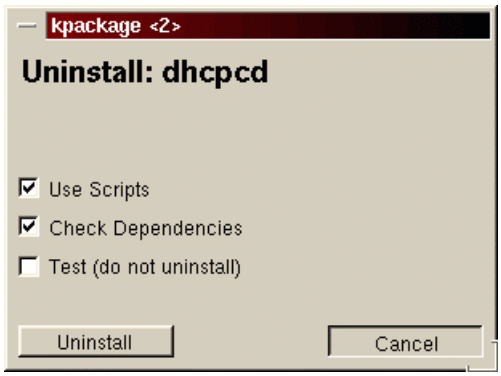

Figure 88. Uninstalling a package

Before you uninstall a package, you can change the options, but we suggest that you leave the default settings unchanged. After you have adjusted the settings, click **Uninstall** to continue. After the dependencies are checked, the package will be uninstalled.

3.5.2 Installing a package

To install a package, click **File > Open**, and you will see a window similar to Figure 89.

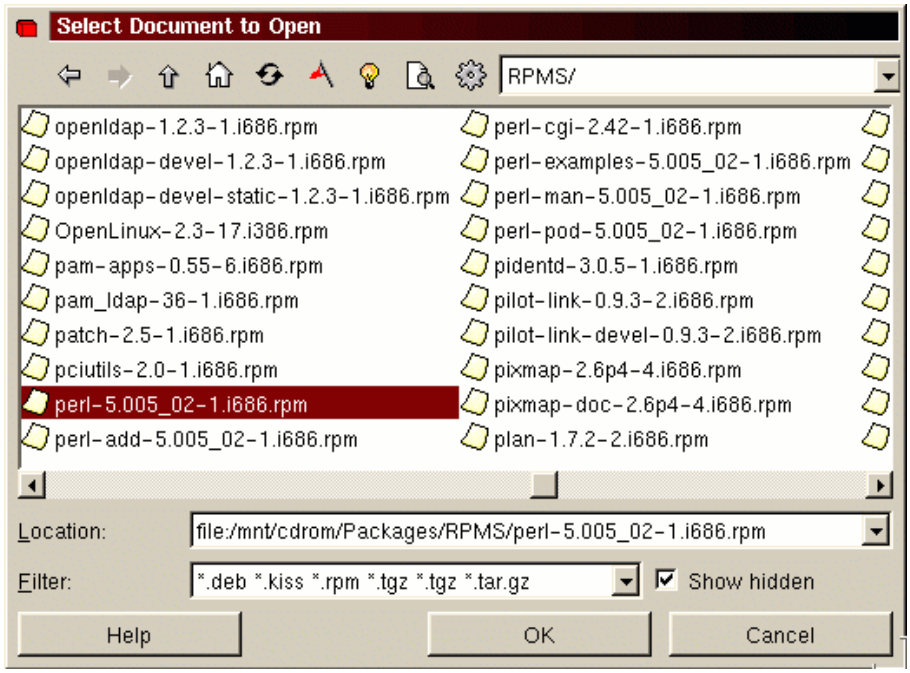

Figure 89. Selecting a package to install

Here you can select the packages you wish to install from any available system directory.

Note

If you want to install packages from a CD-ROM, you must mount the CD-ROM drive before you can access the files on it. This can be done with the command `mount /mnt/cdrom` from a command prompt.

After you have selected the package, click **OK** to continue. You will see a window similar to Figure 90.

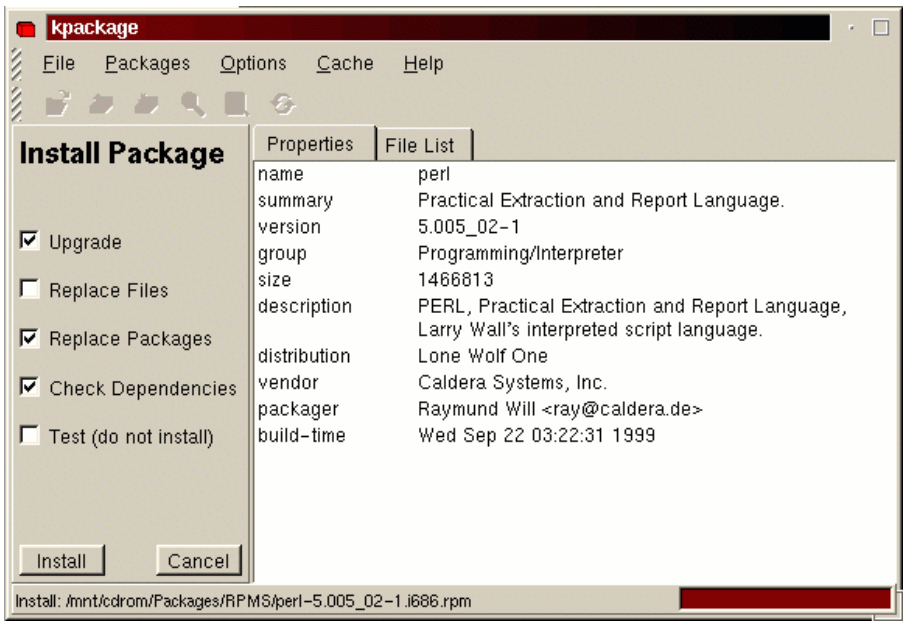

Figure 90. Description of a new package

In this window you see the description of the package. Select the **File List** option to see which files are included with the package. Before you actually install packages, you can adjust the installation options. The options are:

- **Upgrade** - this is used if you are installing a package that is already installed
- **Replace File** - if there are files in the same location already, they will be replaced automatically
- **Replace Packages** - packages are updated in the packages database
- **Check Dependencies** - check if all dependencies are satisfied
- **Test (do not install**) - perform a test installation

After you have selected your options, click **Install** to install the package. After the package is installed it will appear in the package list, as you can see in Figure 91.

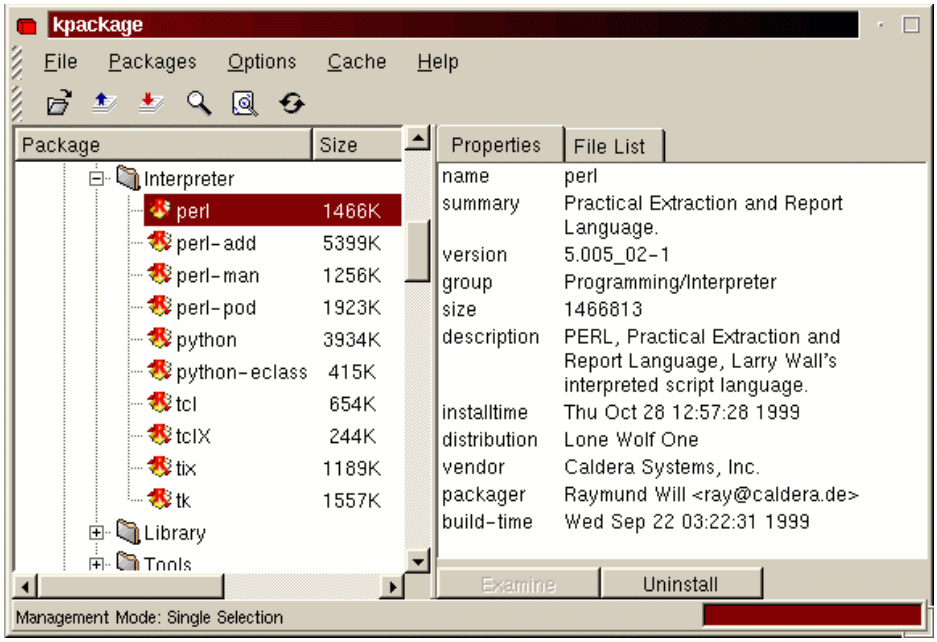

Figure 91. After the installation

3.6 Package management using RPM

Package management can also be done directly with the RPM package manager on the command line. The command line in the graphical interface can be accessed through the terminal window as we described in 3.3, "Getting the X-Windows terminal window" on page 90. Table 9 shows some of the most frequently used versions of the RPM commands.

Table 9. Basic RPM commands

Command	Description
`rpm -q <package>`	If a package is installed, check version and build number of the installed package.
`rpm -qi <package>`	Obtain more information about an installed package.
`rpm -qa`	List all installed packages.
`rpm -qf <filename>`	Determine the (installed) package that `<file>` belongs to.

Command	Description
`rpm -Uhv <package.rpm>`	Update/Install the file `package.rpm` showing a progress bar.
`rpm -F -v ./*.rpm`	Update (freshen) all currently installed packages using the RPM files in the current directory.
`rpm --help`	Get help about the different options and parameters.

Note

After you install packages using RPM, you may need to run some additional configuration programs. Programs such as Apache need to be customized to your particular environment and require some post-installation maintenance. Some of these packages can be configured from the graphical interface by selecting other icons. Other packages have their own configuration tools.

More information and options about RPM can be found in the manual page (man rpm), the RPM how-to file (less /usr/doc/howto/en/RPM-HOWTO.txt.gz) and at the RPM Web site at `http://www.rpm.org`. You can also display a short overview by running `rpm --help`..

3.7 System menu

In the System menu of the COAS tools, you can access the following tools:

- Accounts - for managing the accounts
- Daemons - for managing the startup programs
- Filesystem - for mounting devices and NFS volumes
- Hostname - for setting hostnames
- Resources - for checking the hardware resources
- Time - for setting the time and time zone

The System menu is shown in Figure 92.

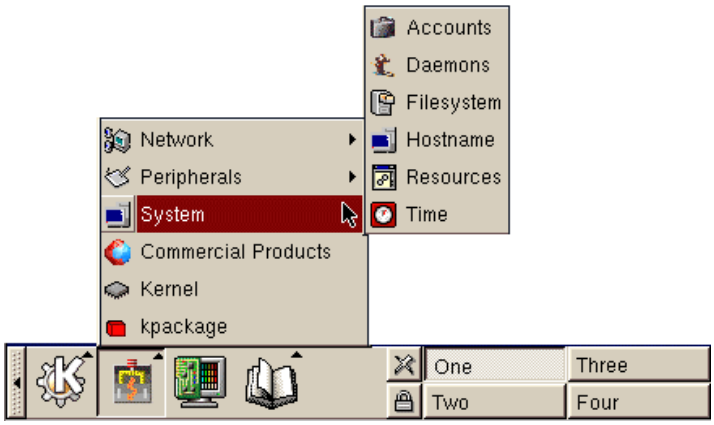

Figure 92. System menu

To start the tools from the System menu, select the tool you want. At the initial window, click **OK** to continue.

3.8 Accounts

This tool is used to manipulate the user accounts. After the tool is started, you will see a window similar to Figure 93.

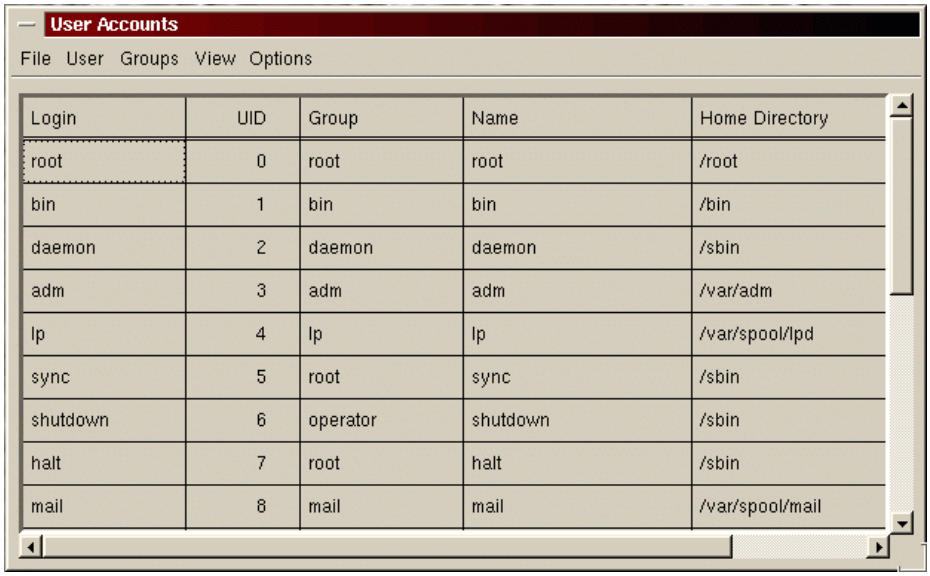

Figure 93. Account management

Here you can manage users and groups. In the following sections we will describe how to perform these tasks.

In the View menu, you have two options for displaying users:

- All users - all users will be displayed
- Regular users - only regular users will be displayed

In the Options menu, you have three options to choose from:

- Preferences - here you define the global preferences for creating users and groups. If you select this option you will see a window similar to Figure 94.

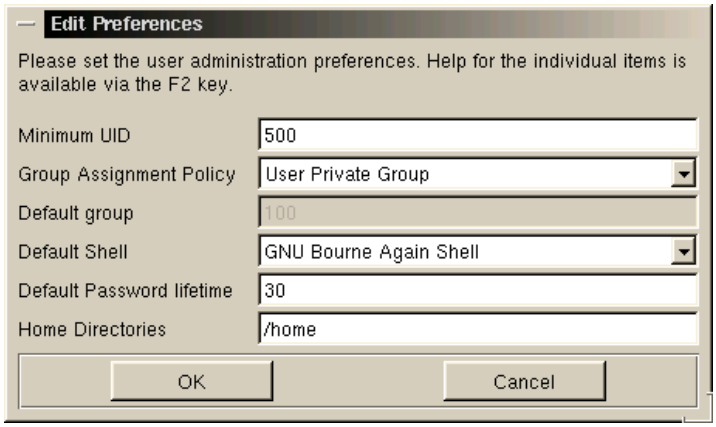

Figure 94. Setting the preferences for creation

Define your preferences and click **OK** to store them.

- Enable/Disable shadow passwords - here you can enable or disable shadow passwords.
- Enable/Disable NIS lookups - here you can enable or disable NIS lookups.

3.8.1 Managing accounts

In this section we explain how to manage accounts. We cover adding a new user, deleting a user and editing an existing user.

To create a new user follow these steps:

1. To add a new user, select **User > Create User**. You will see a window similar to Figure 95.

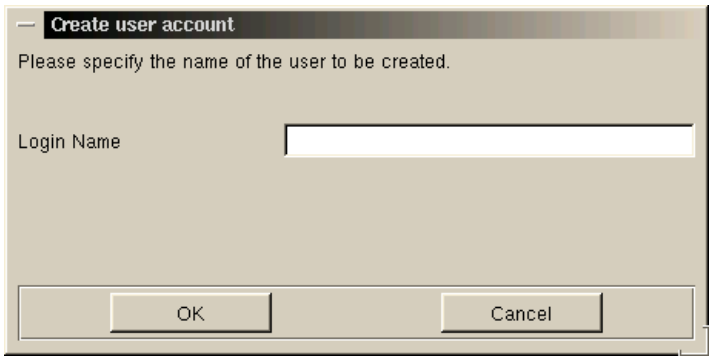

Figure 95. Login name for the new user

2. Type in the unique login name of the new user and click **OK** to continue, and you will see a window similar to Figure 96.

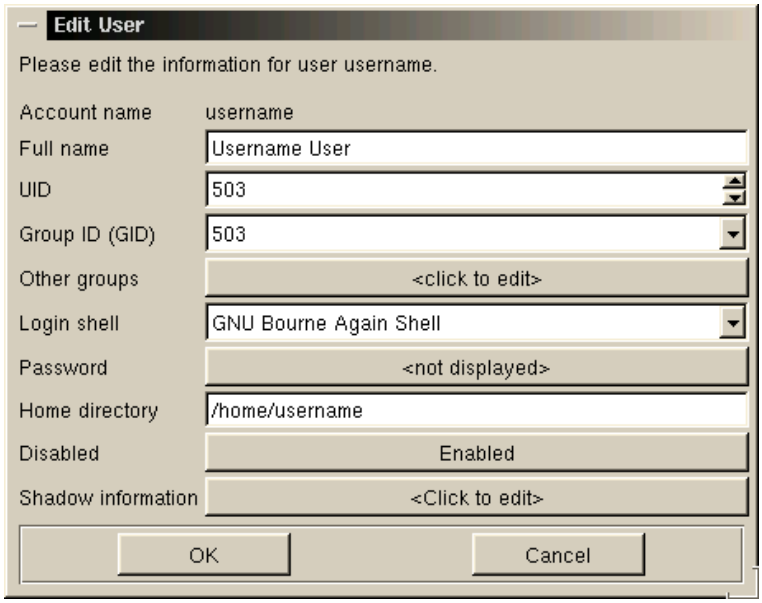

Figure 96. Specifying parameters for the new user

Here you need to specify the following:

- **Full name** - this is the description of the user
- **UID** - this is the number by which the system knows you. It only attaches this number to file and directory ownership and uses /etc/passwd to convert this to a username when listing the attributes. Generally UID numbers are unique and the system programs will usually prevent you from creating more than one username with the same UID. This can usually be overridden by specifying options to the commands to create IDs.
- **GID** - this is a unique number assigned to a group. In Caldera OpenLinux each user has its own default group. The default GID is the next available and the GID numbers are starting at 500.
- **Other groups** - each user can be a member of one or more groups. You can specify these groups here. If you click the button you will see a window similar to Figure 97.

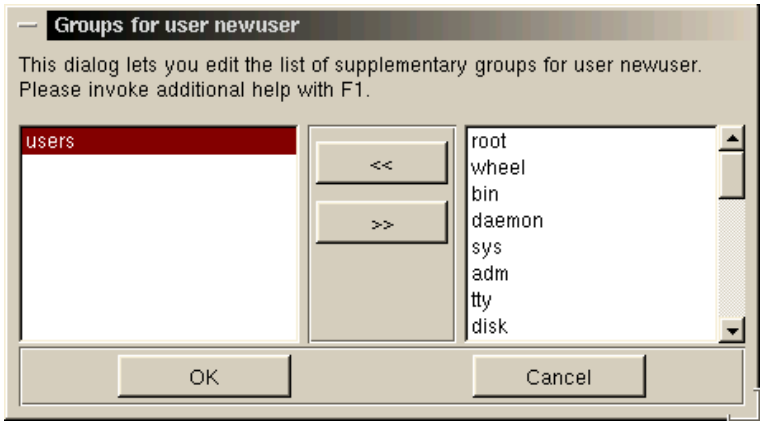

Figure 97. Specifying other groups for the user

When you have added all the groups you want, click **OK** to continue.

- **Login shell** - the shell that is started when the user logs in.
- **Password** - the password used to log in with. To define a password, click the button labeled **<not displayed>** and you will see a window similar to Figure 98.

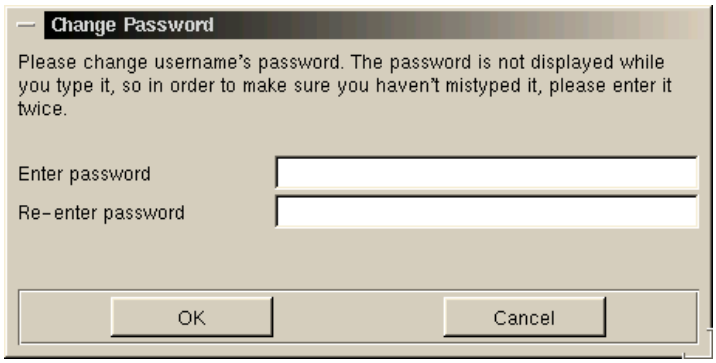

Figure 98. Specifying the password for the new user

Type in the password for the user and click **OK** to continue.

Note

Caldera OpenLinux uses shadow passwords by default.

- **Home directory** - this is the user's home directory. It is the first place a user goes to when logging in. It contains files and programs that are owned and used by that user.
- **Disabled/Enabled** - with this you define if an account is enabled or disabled. You can toggle this value by clicking the button.
- **Shadow information** - here you define the password properties: expiration, change timeframe, etc. If you want to change the default values, click the button and you will see a window similar to Figure 99.

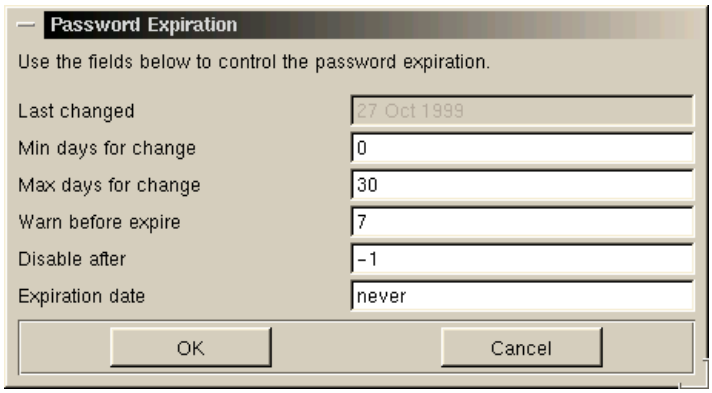

Figure 99. Password properties for the new user

When you have edited the properties, click **OK** to save them.

3. After you have typed in all necessary information for the new user, click **OK** to actually create the new user.

3.8.1.1 Deleting a user

When you want to delete a user, select the user from the list and click **User > Delete User**. You will see a window similar to Figure 100.

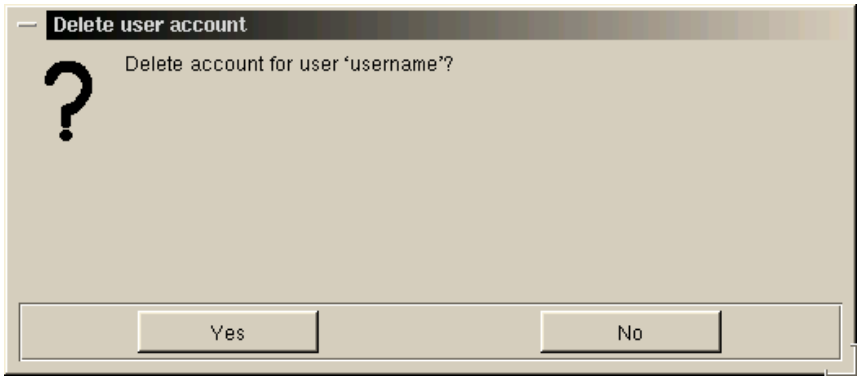

Figure 100. Deleting a user

Click **Yes** to delete the user.

3.8.1.2 Editing a user

When you want to edit a user, select the user from the list and choose **User > Edit User**. You will see a window similar to Figure 101.

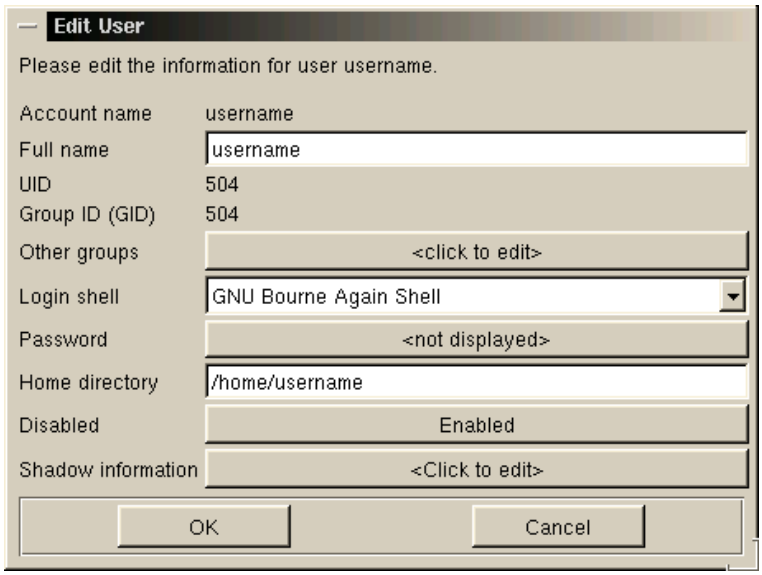

Figure 101. Edit a user

Here you can modify the following attributes of a user:

- Other groups
- Login shell

- Password
- Home directory
- Disabled/Enabled
- Shadow information

We described these attributes in 3.8.1, “Managing accounts” on page 101. When you are done, click **OK**.

3.8.2 Managing groups

You can access the tool for managing groups by selecting **Accounts > Groups > Manage groups**. You will see a window similar to Figure 102.

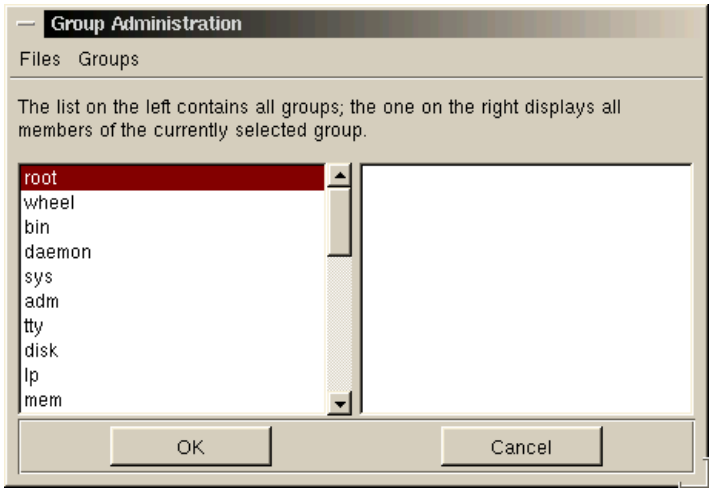

Figure 102. Group Administration

Here you can perform operations related to the groups.

3.8.2.1 Creating a new group

You can create a new group by selecting **Groups > Create Group**. You will see a window similar to Figure 103.

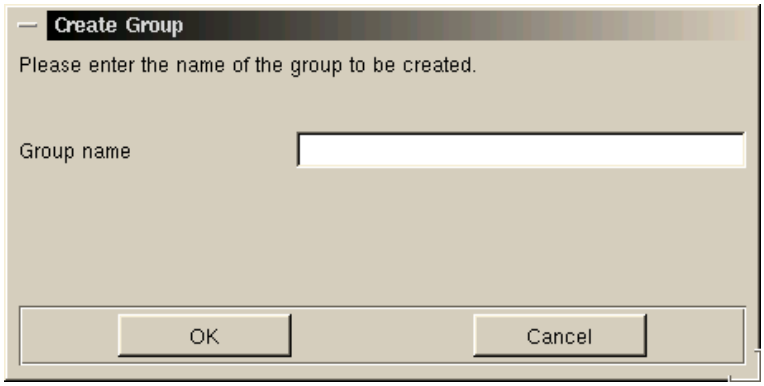

Figure 103. Creating a new group

Type in the name of the new group and click **OK** to create it.

3.8.2.2 Deleting a group

Select the group you want to delete from the list of all the groups and choose **Groups > Delete Group**. You will see a window similar to Figure 104.

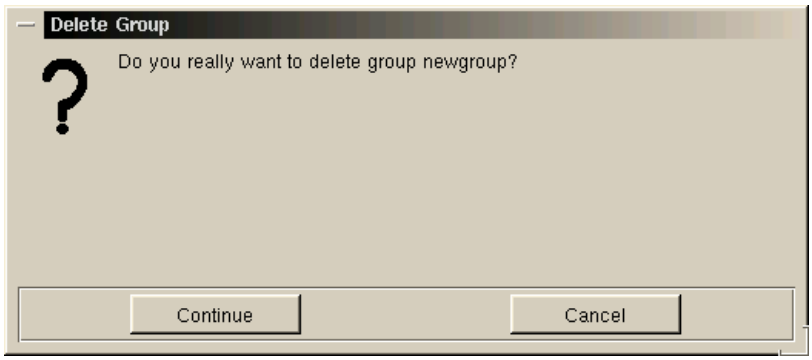

Figure 104. Deleting a group

Click **Continue** to actually delete a group.

3.8.2.3 Rename a group

Select the group you want to rename from the list of all the groups and choose **Groups > Rename Group.** You will see a window similar to Figure 105.

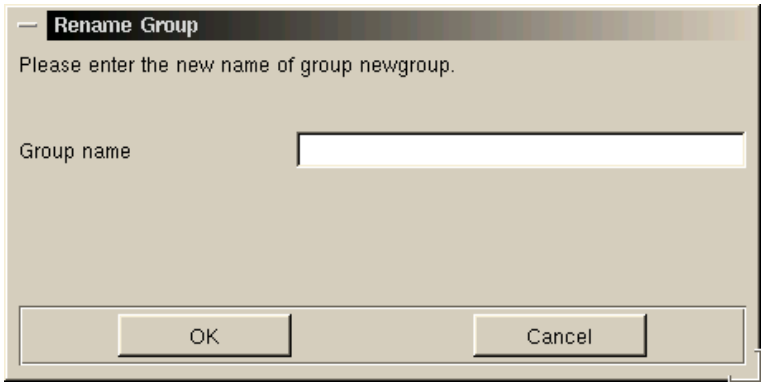

Figure 105. Renaming a group

Type in the new name for the group and click **OK** to rename it.

3.8.2.4 Merge a group

You have the option to merge users from one group to another. Select the group to which you want to merge another and choose **Groups > Merge Group**. You will see a window similar to Figure 106.

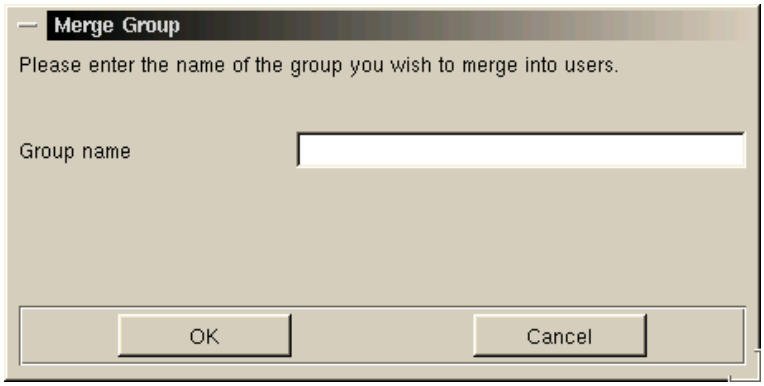

Figure 106. Merge a group

Type in the name of the group you want to merge in. Click **OK** to continue.

3.8.2.5 Group membership

You can change the members of a group. To change the members of a desired group select the group from the list of all the groups and choose **Groups > Group Membership**. You will see a window similar to Figure 107.

Figure 107. Group membership

You can add or remove users from a group. Click **OK** to save your changes.

3.9 Daemons (services)

This tool is used to manipulate the daemons that will start at the server startup. After the tool is started you will see a window similar to Figure 108.

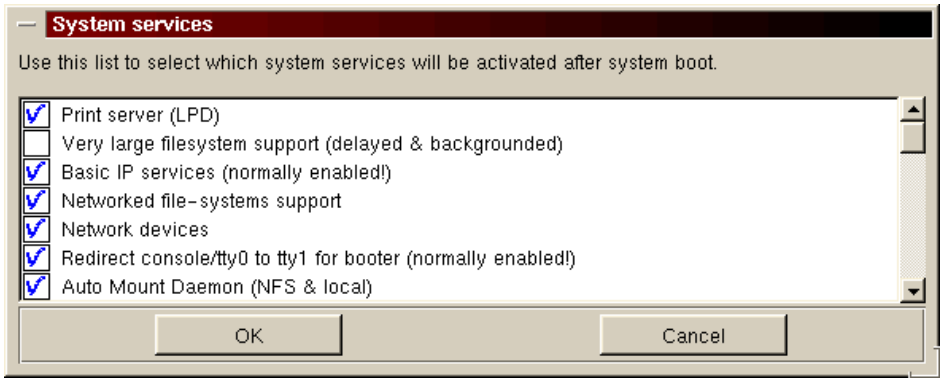

Figure 108. System services

Here you define which services (daemons) will be started at the server startup. When you are finished, click **OK** to save your changes.

3.10 Filesystem

Here you can mount or unmount the devices and connect to the NFS servers. After the tool is started, you will see a window similar to Figure 109.

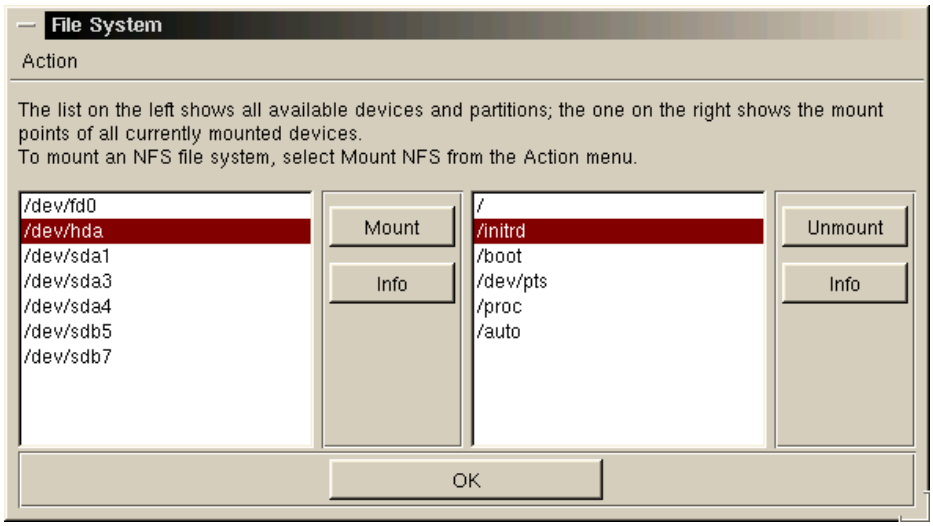

Figure 109. Filesystems

On the left side you see unmounted devices. If you want to mount the device, select it from the list and click **Mount.**

On the right side you see mounted devices. If you want to unmount an already mounted device, select it from the list and click **Unmount.**

By selecting the mounted or unmounted device and clicking **Info,** you will see the information about the particular device.

3.10.1 Mounting an NFS volume

You can mount an NFS file system by choosing **Action > Mount NFS**. You will see a window similar to Figure 110.

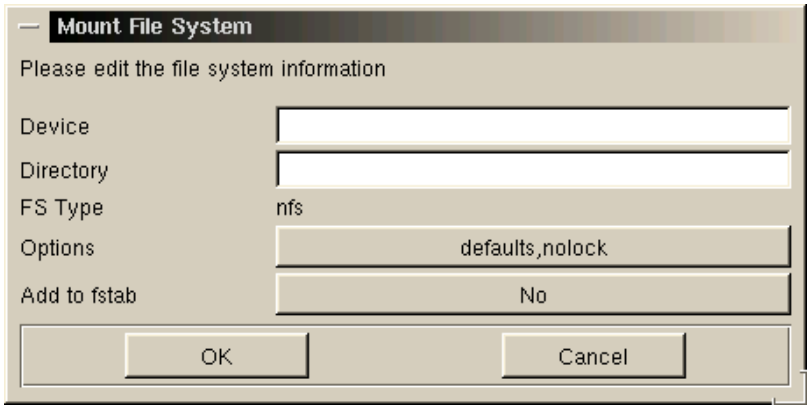

Figure 110. Mounting an NFS volume

Type in the required values and click **OK** to mount the NFS volume.

3.11 Hostname

Here you can change the hostname of your Linux server. After the tool is started you will see a window similar to Figure 111.

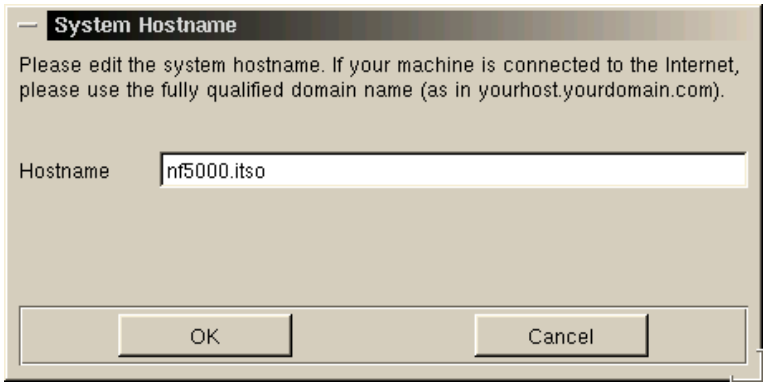

Figure 111. Changing the hostname

Type in the new hostname and click **OK** to save it.

3.12 Resources

With this tool you can examine hardware resources. After the tool is started you will see a window similar to Figure 112.

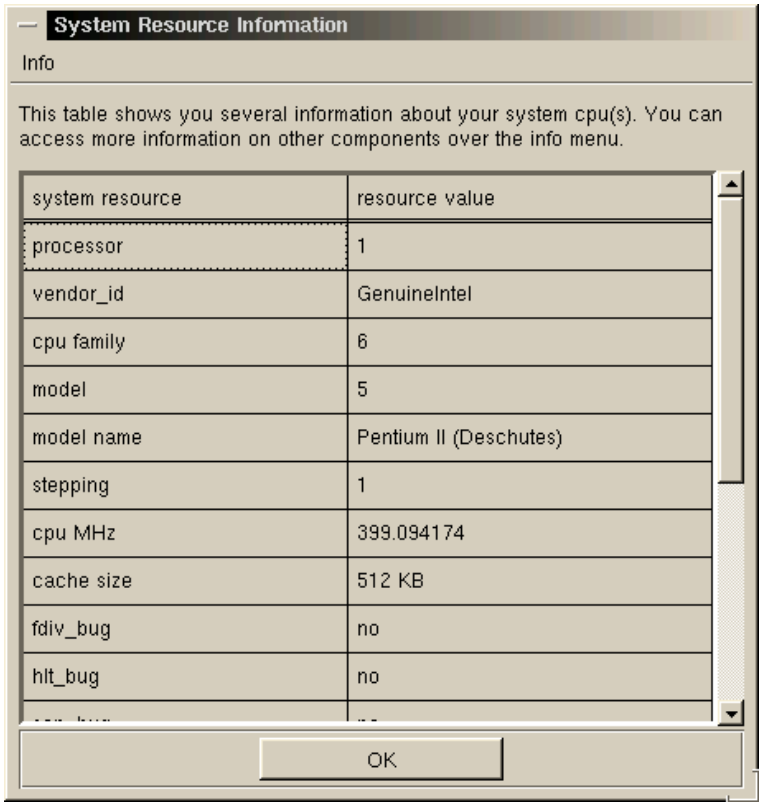

Figure 112. System resources

Here you can get information about the following resources:

- Block devices
- Character devices
- Interrupts
- System load average
- IOports
- DMA

To access this information, select the appropriate option from the Info menu. For example if you select **Interrupts**, you will see a window similar to Figure 113.

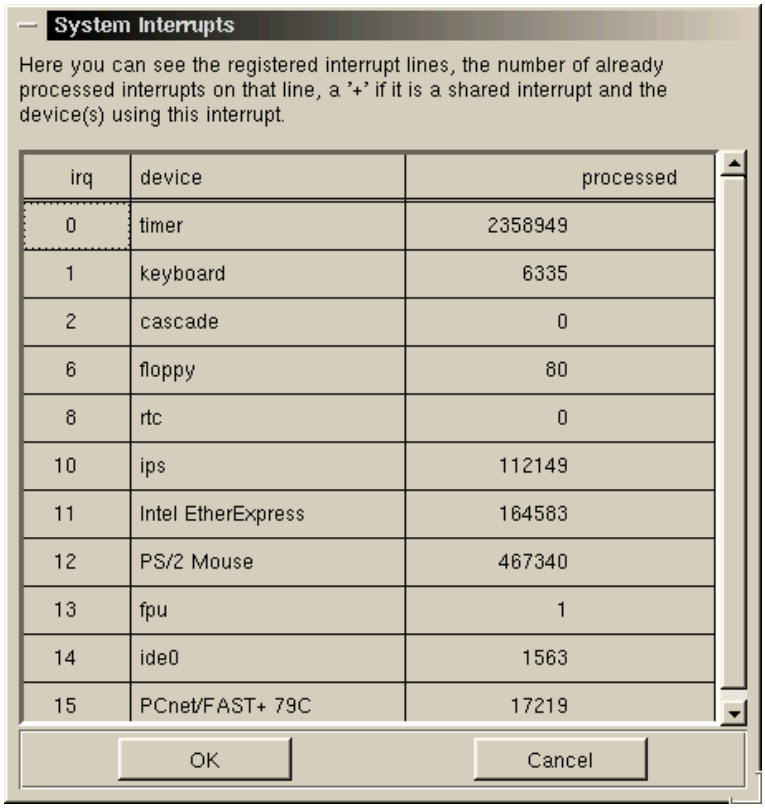

Figure 113. System interrupts

3.13 Time

Here you can set the time and time zone. After the tool is started, you will see a window similar to Figure 114.

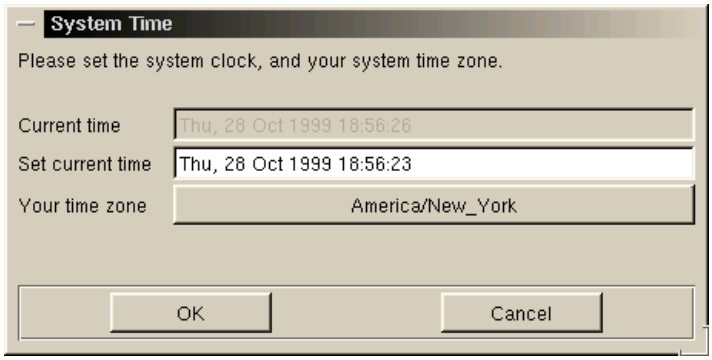

Figure 114. Setting the time

Type in the current time. If you also want to change the time zone, click the button for it. You will see a window similar to Figure 115.

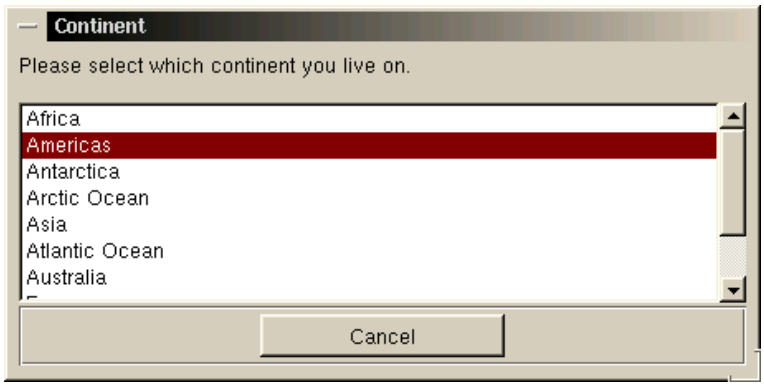

Figure 115. Setting the time zone

Select your region and you will be presented with the time zones for that region. Select the one that matches your city. After that you will be back in the System Time panel. Click **OK** to save the changes.

3.14 Peripherals menu

In the Peripherals menu of the COAS tools you can access the following tools:

- **Mouse** - for managing the mouse
- **Printers** - for managing the printers

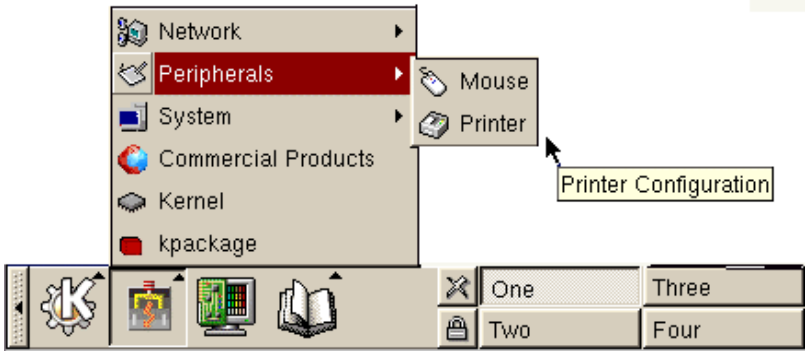

Figure 116. Peripherals menu

To start the tool from Peripherals menu, select the tool you want. At the initial dialog, click **OK** to continue.

3.15 Mouse

This tool is used to configure the behavior of the mouse in the text-based user interface. After the tool is started you will see a window similar to Figure 117.

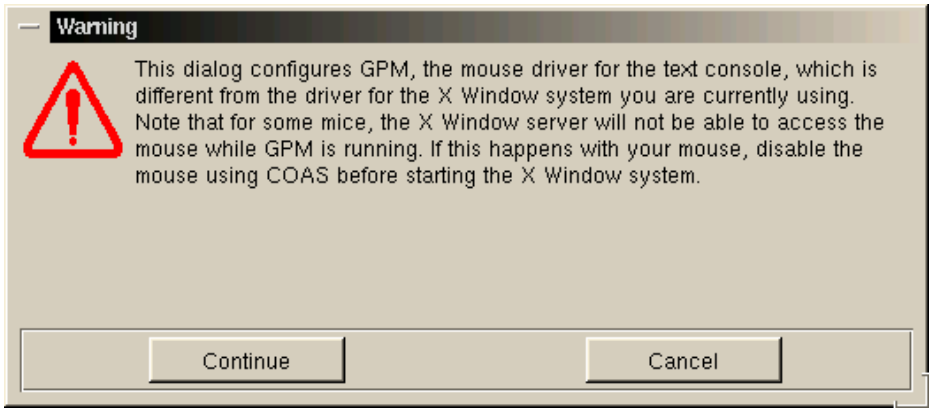

Figure 117. Warning before configuring the mouse

As you can see from the warning, this tool is used for configuring the GPM to enable additional features for mouse usage in the text-based interface. Click **Continue** to continue with the configuration. You will see a window similar to Figure 118.

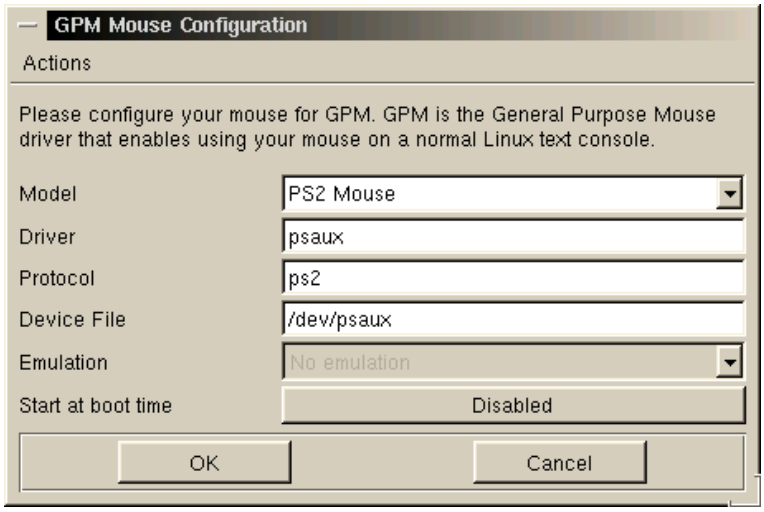

Figure 118. GPM Mouse Configuration

Select the configuration parameters that meet your needs and click **OK** to continue. On the next window, click **Save** to save your settings.

Note

If you did not install the GPM package, you will receive the error message that the daemon cannot be started.

3.16 Printer

This tool is used to configure the printers you want to use in your Caldera OpenLinux system. After the tool is started, you will see a window similar to Figure 119.

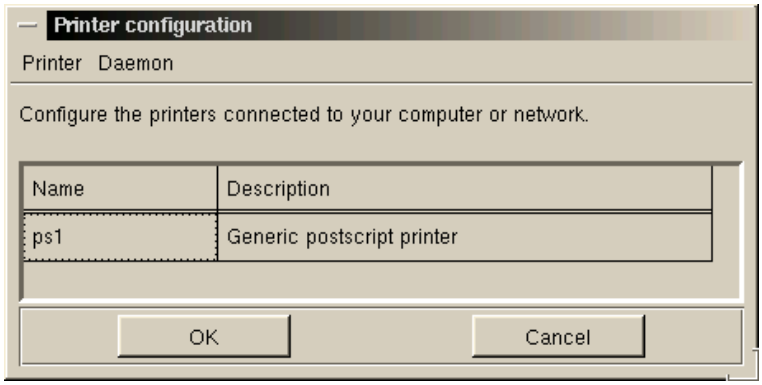

Figure 119. Printer configuration

Here you can manage printers. In the following sections we will describe how to perform these tasks.

From the Daemon menu you can start or stop the printer daemon.

Note

You can only print documents if the daemon is running.

3.16.1 Adding a new printer

You can add a new printer to your system by selecting **Printer > Add**. You will see a window similar to Figure 120.

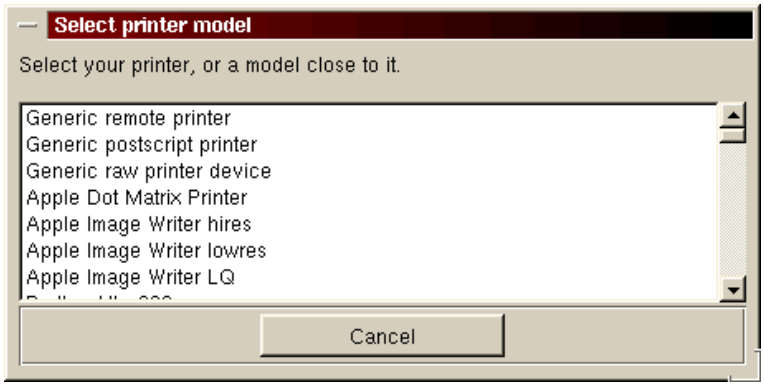

Figure 120. Selecting a printer model

Select your model from the list. After that you will see a window similar to Figure 121.

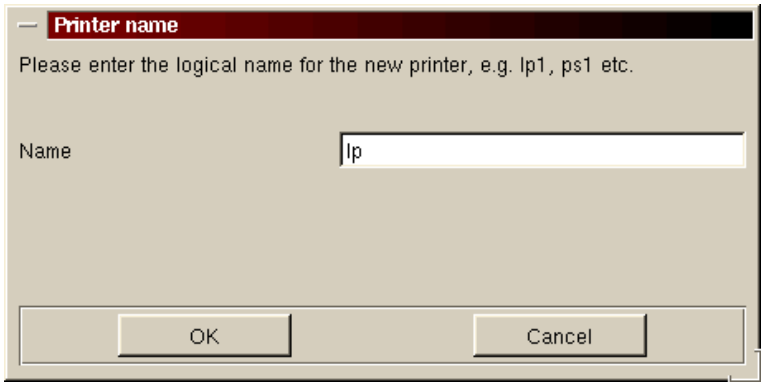

Figure 121. Defining printer logical name

Here you define the logical name of the printer. This name is then used in all printer definitions. Click **OK** to continue, and you will see a window similar to Figure 122.

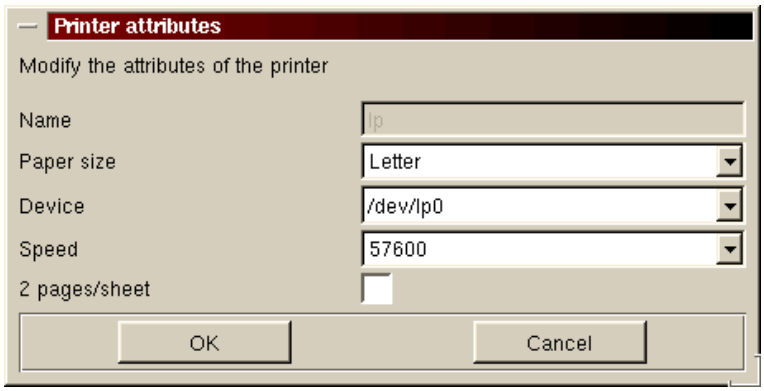

Figure 122. Printer attributes

Here you define printer attributes:

- **Paper size**
- **Device** - this is the physical device to which the printer is connected. This is usually the parallel port, and /dev/lp0 is the first parallel port in your server.
- **Speed** - this is the speed for the data traveling over the device to which the printer is connected.

These attributes are related to the printer driver you choose, so all drivers will not have the same options.

After you have defined all attributes for your printer driver click **OK** to continue. In the next window select **Save** to save your configuration. The installation program will then ask you if it should create the printer queue for your new printer. Click **OK** to create the queue. The printer daemon will be stopped so that the queue can be created and then it will be restarted again.

3.16.2 Removing a printer

You can remove a printer from your system by selecting the printer to be removed from the list of installed printers and select **Printer > Remove**. You will be asked twice if you really want to remove the printer. Answer **OK** both times if you really want to remove the printer.

3.16.3 Edit a printer

If you want to edit the properties of the installed printer, select the printer from the list and choose **Printer > Edit**. You will see a window similar to Figure 123.

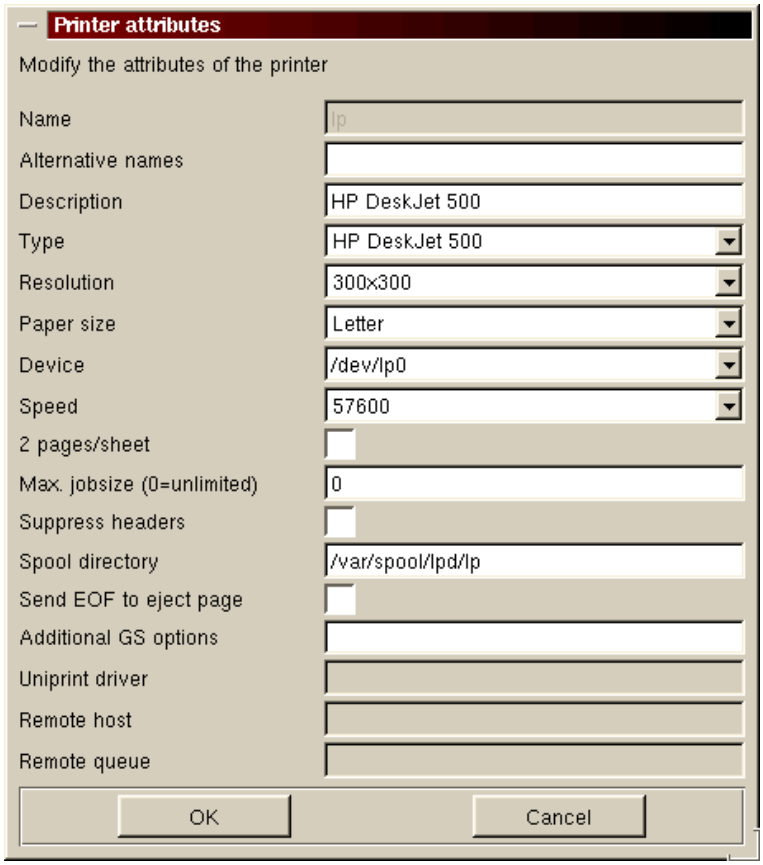

Figure 123. Printer attributes

Edit the preferences you want and click **OK** to continue. In the next window click **Save** to save the changes.

3.17 Network menu

From the Network menu of the COAS tools, you can access the following tools:

- **TCP/IP** - for managing TCP/IP settings.
- **Ethernet interfaces** - for Ethernet Network Interface Cards (NICs).
- **Mail Transfer** - for managing the Mail Transfer Agent (MTA). You can find more information on how to set up MTA on your server in Chapter 9, “sendmail” on page 213.

The COAS tools Network menu is shown in Figure 124.

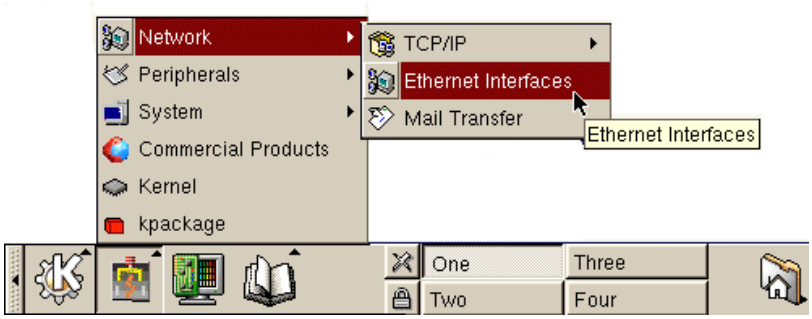

Figure 124. Network menu

To start the tools from the Network menu select the tool you want. At the initial window, click **OK** to continue. If you select **TCP/IP,** you will be presented with two options, as you can see in Figure 125.

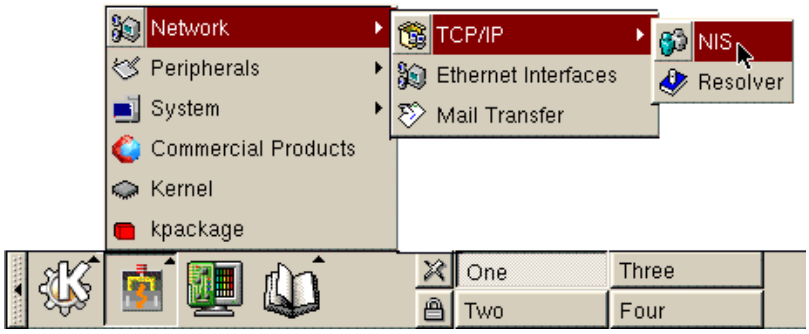

Figure 125. TCP/IP menu

- NIS - for setting the NIS client options. You can get more information on how to set up an NIS client or server in Chapter 11, “NIS - Network Information System” on page 247.
- Resolver - to set up the TCP/IP resolving settings.

3.18 Ethernet interfaces

With this tool you can configure your Ethernet NICs. After you start the tool you will see a window similar to Figure 126.

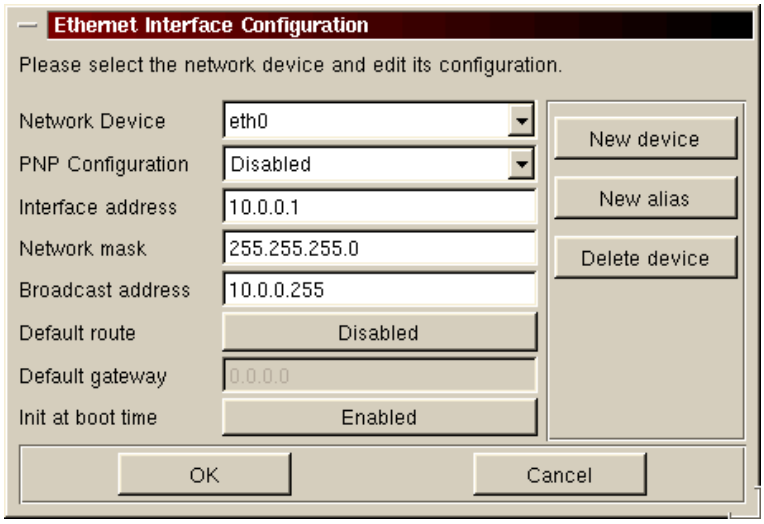

Figure 126. Ethernet Interface Configuration window

If you configured your Ethernet NIC during installation you will see the current configuration. There are several configuration options available:

- Network device - this is the name of the network device as it is recognized by the kernel.
- PNP Configuration - here you can select if the adapter is configured automatically from a DHCP server by selecting the **DHCP** option, or manually by selecting the **Disabled** option.
- Interface address - here you define the IP address of the interface.
- Network mask - here you define the subnet mask for the interface.
- Broadcast address - here you define the broadcast IP address. This is by default calculated from subnet mask.
- Default route - here you enable or disable the default route.
- Default gateway - if you enabled default routing, you need to specify the IP address of the router here.
- Init at boot time - here you specify if the interface should be initialized during system startup.

3.18.1 Adding a new network interface

If you have installed a new Ethernet interface you can add it to the system configuration by clicking **New device.** You will see a window similar to Figure 127.

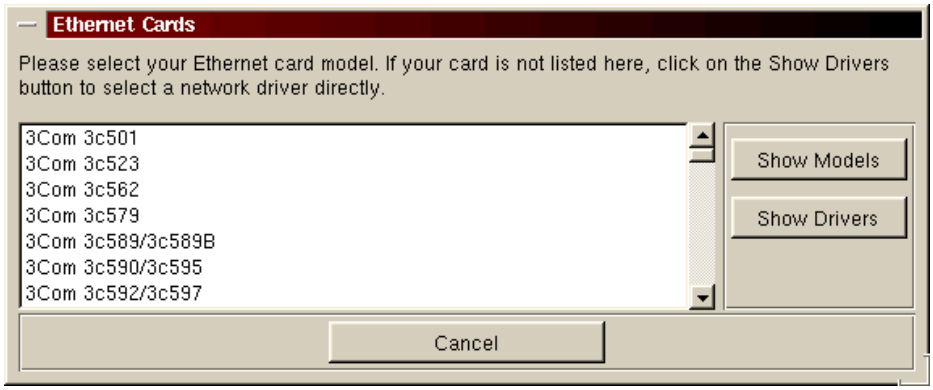

Figure 127. Selecting the type of the Ethernet card

If you do not find the driver for your Ethernet card among the listed models, you may try to check all available drivers. To see all drivers click **Show Drivers.** Select your model/driver by clicking the appropriate one, and you will see a window similar to Figure 128.

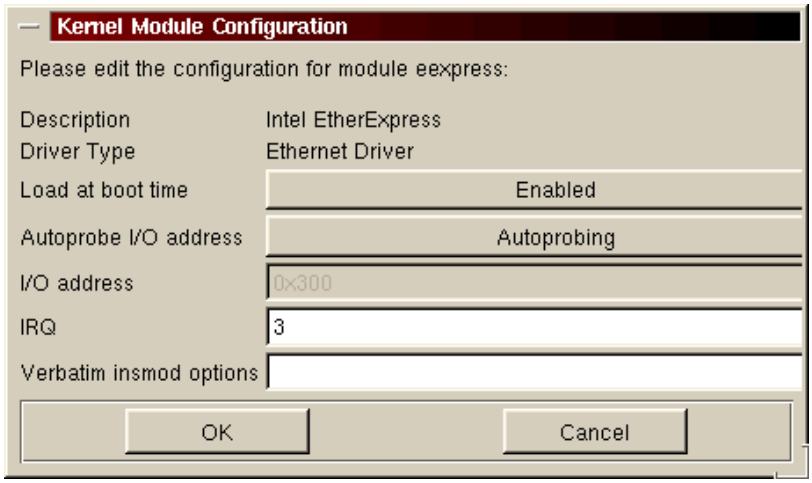

Figure 128. Defining hardware parameters

Here you define the hardware parameters for the driver for your Ethernet NIC. When you are done, click **OK** to continue. The setup utility will try to load the module you selected. If the loading of the module is successful, your new interface definition will now be available for additional setup. You will see a window similar to Figure 126 on page 122. Define parameters to meet your needs and click **OK** to continue. In the next window click **Save** to save the configuration.

3.18.2 Removing a network interface

If you want to delete the definition for an Ethernet NIC, click **Delete device** from the dialog shown in Figure 126 on page 122. Then click **OK** to close the configuration window. In the next window, click **Save** to save the changes you just made.

If you have more than one Ethernet NIC adapter and you want to remove the adapter eth1 for example, follow these steps:

1. Stop the interface by executing the command:

   ```
   /sbin/ifdown eth1
   ```

2. Delete the file /etc/sysconfig/network-scripts/ifcfg-eth1 by executing the command:

   ```
   rm /etc/sysconfig/network-scripts/ipcfg-eth1
   ```

This procedure can be used for all adapters when you have multiple adapters defined.

3.19 Name resolution settings

You can access the tool for name resolution settings by clicking **Network > TCP/IP > Resolver**. When you start the tool you will see a window similar to Figure 129.

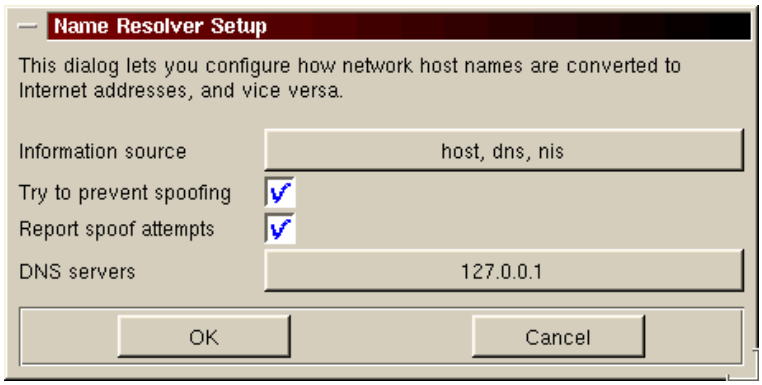

Figure 129. Name resolution setup

Here you can define how the name resolution is performed on your system. You have four options here:

- Information source - here you define the order and sources for the name resolution.

- Try to prevent spoofing
- Report spoof attempts
- DNS servers - the defined IP addresses of the DNS servers

3.19.1 Name resolution order and sources

You can change the name resolution order and sources by clicking the button to the right of **Information sources.** You will see a window similar to Figure 130.

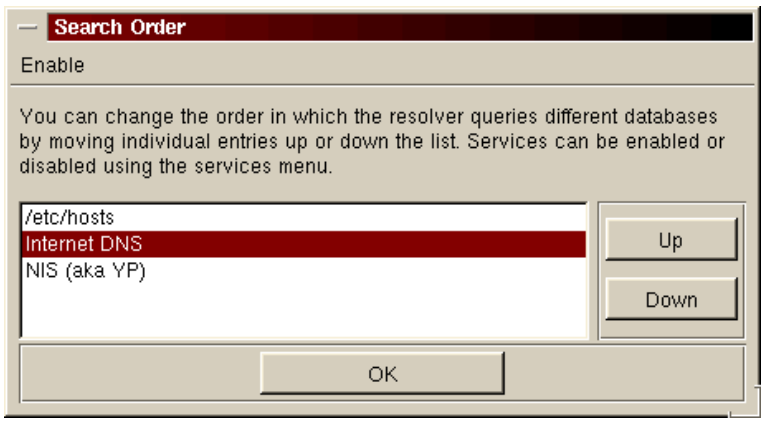

Figure 130. Search order

The search order can be changed by moving the name resolution resources up and down. If you want to enable or disable a particular name resolution source you can do this by selecting **Enable** and selecting the source you want to enable or disable. If a source is currently enabled, you can disable it and vice versa. When you are done, click **OK** to continue and on the next window select **Save** to save the changes.

3.19.2 Defining a DNS server

You can define a DNS server by clicking the button to the right of the **DNS servers** button. You will see a window similar to Figure 131.

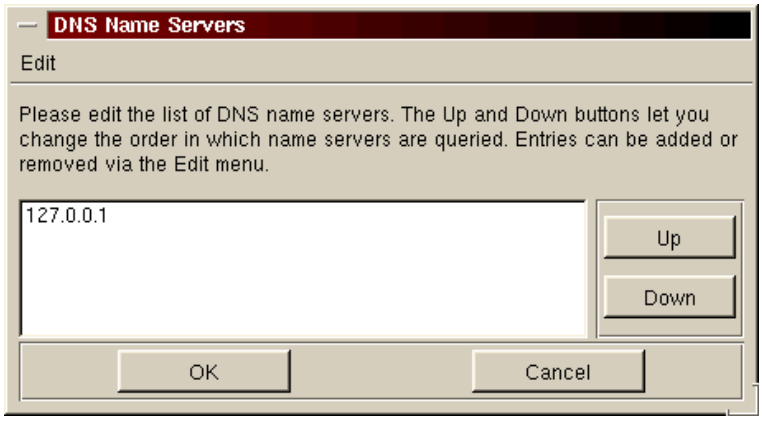

Figure 131. DNS servers

If you have more than one DNS server defined you can reorder them by moving them up and down. The top-most server will be accessed first and so on.

3.19.2.1 Add a new DNS server

If you want to add a new DNS server select **Edit > Add server**. You will see a window similar to Figure 132.

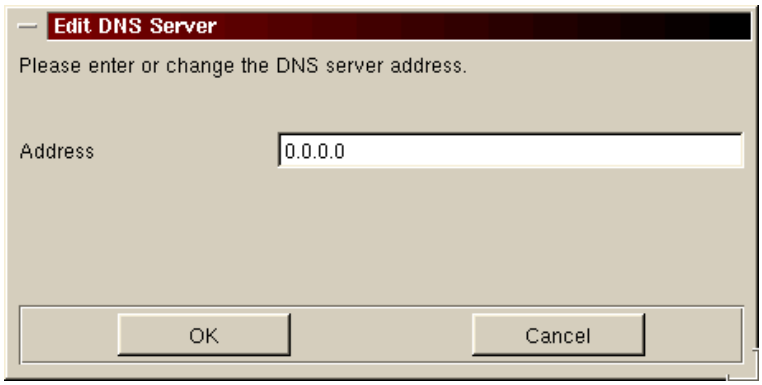

Figure 132. Specifying a DNS server

Type in the IP address of the DNS server and click **OK** to go back to the previous window.

3.19.2.2 Remove a DNS server

If you want to remove a DNS server select it from the list and choose **Edit > Remove server**.

3.19.2.3 Change a DNS server

If you want to change a DNS server's IP address select it from the list and choose **Edit > Edit server**. You will see a window similar to Figure 132. Type in the new IP address of the server and click **OK** to go back to the previous window.

3.20 Manipulating kernel modules

You can manage kernel modules in Caldera OpenLinux by using the kernel configuration tool from the COAS tools. You can start it by selecting **Kernel** from the COAS tools menu. When the Kernel tool is started, you will see a window similar to Figure 133.

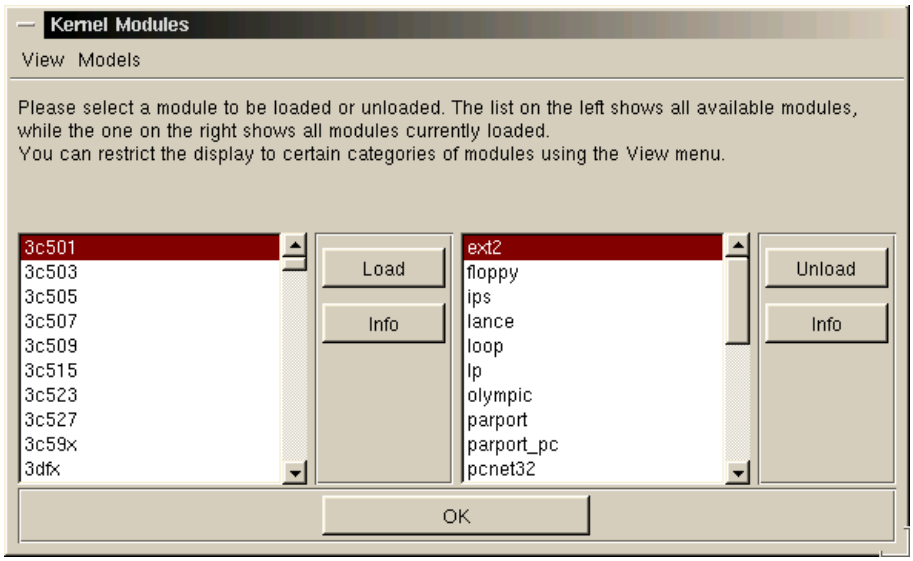

Figure 133. Kernel Modules

On the left side you can see all available modules and on the right side you can see loaded modules. By default all modules are displayed, but if you want to display just one kind of module, you can do this by selecting the following options from the View menu:

- All drivers
- Arcnet drivers
- CD-ROM drivers
- Ethernet drivers
- Misc drivers
- Network drivers

- SCSI drivers
- SCSI host adapter drivers
- Sound drivers
- Token-ring drivers
- ISDN drivers
- Multimedia drivers

If you want to get information about a particular module, select the module from either side and click **Info.**

3.20.1 Loading a new module

When you install a new piece of hardware you need to load the appropriate module if you want the hardware to be useful. In Linux, drivers can be loaded or unloaded without restarting the system. It may take some time to get used to this if you are used to another popular operating system. To load a new module select the module from the left side and click **Load.** You will see a window similar to Figure 134.

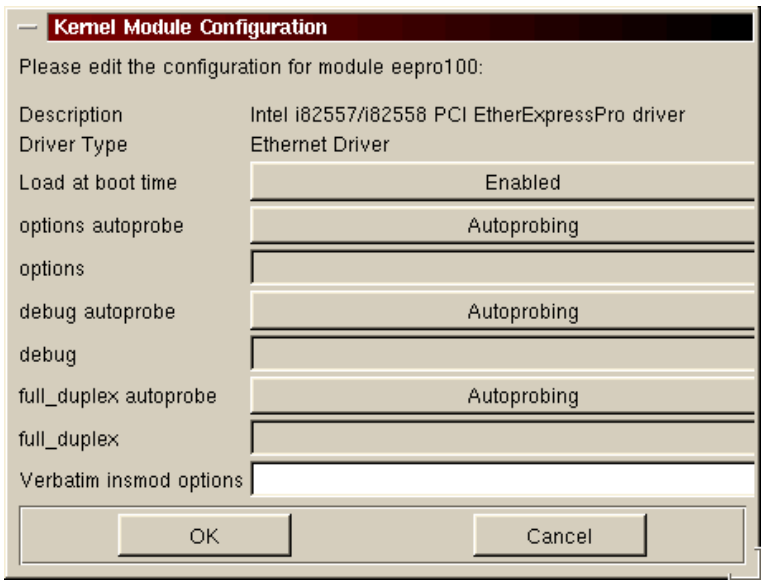

Figure 134. Module configuration

Each module has several hardware-related options and an option to load at boot time. If you want to load a module at boot time, click the button to the right of the Load at boot time field to specify your preferred setting. Click **OK** to actually load the module. If the module is loaded successfully, it will appear on the left side where the loaded modules are displayed.

3.20.2 Unloading a new module

If you want to unload an already loaded module, select it from the left side and click **Unload.** You will be asked if you really want to unload the module. Click **OK** to unload the module. If the module has been enabled to load at system startup, you will see a window similar to Figure 135.

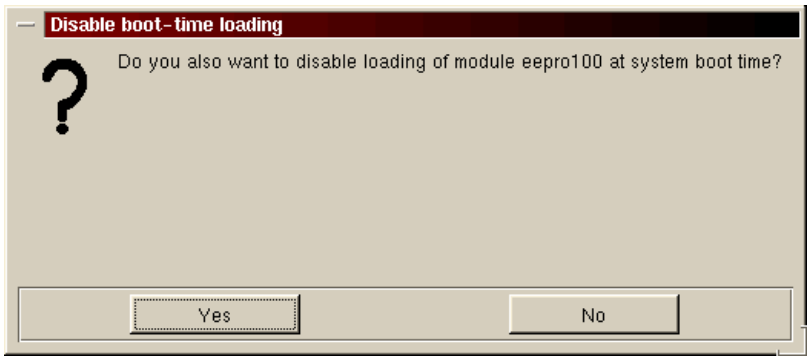

Figure 135. Disabling loading at startup

Here you can decide if you will also disable the startup loading of the module. Select **Yes** or **No** to continue. After that the module will be unloaded.

3.21 Configuring X-Windows

If for whatever reason you need to change the X-Windows setup after installation, you can by executing XF86Setup from the command line.

3.22 System administration using Webmin

In Caldera OpenLinux eServer 2.3, you can also perform administration with the Webmin tool. This is a Web-based interface for managing. Webmin is basically an HTTP daemon acting as an interface to the system files.

Before starting using it, install the latest available version from:

```
http://www.webmin.com
```

You can start it by pointing the Web browser to the IP address of the server on port 1000 as you can see in Figure 136.

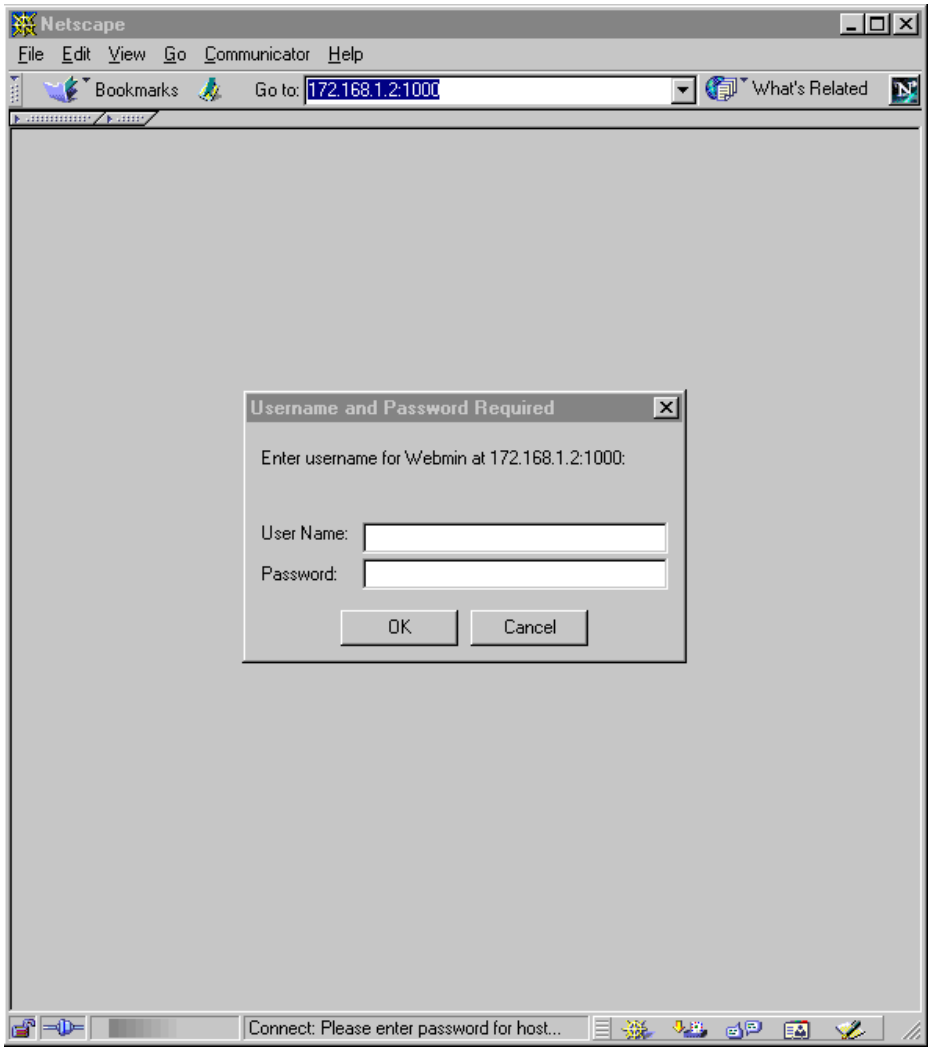

Figure 136. Starting Webmin

Note

In latest Webmin version the access port is by default 10000.

To access the services you need to log on with the **root** user. After logging on you will see a window similar to Figure 137.

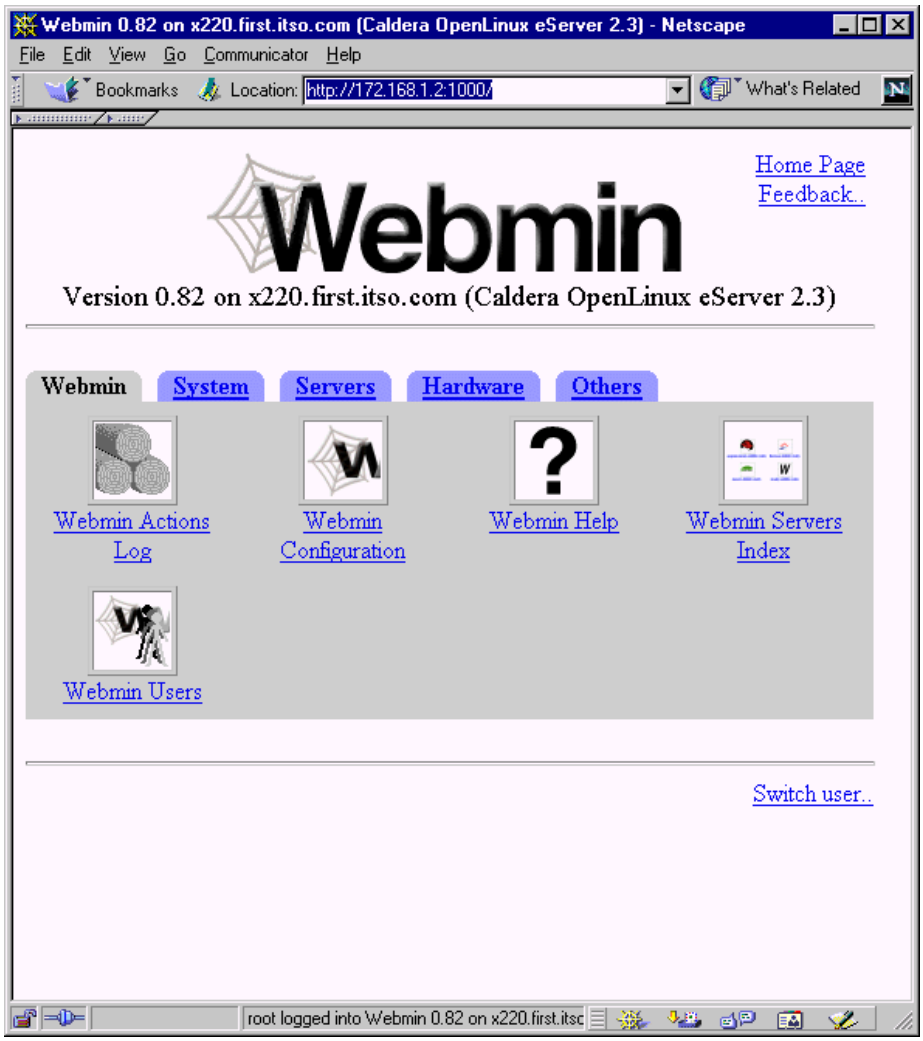

Figure 137. Webmin main window

As you can see the main window is divided into five sections:

- Webmin
- System
- Servers
- Hardware
- Others

In the following sections we will briefly describe the possible configuration operations in each of this sections.

3.22.1 Webmin section

As you can see in Figure 137 on page 131, on the Webmin tab you have the following configuration tasks available:

- Webmin Actions Log - here you can search the Webmin logs
- Webmin Configuration - here you can configure the setup of the Webmin
- Webmin Help - here you can get the help on Webmin modules
- Webmin Servers - here you can search for other Webmin server on the network
- Webmin Users - here you can create, modify or delete users for accessing the Webmin

3.22.2 Webmin system

After you click the **System** tab in the Webmin main window similar to Figure 137 on page 131, you will see a window similar to Figure 138.

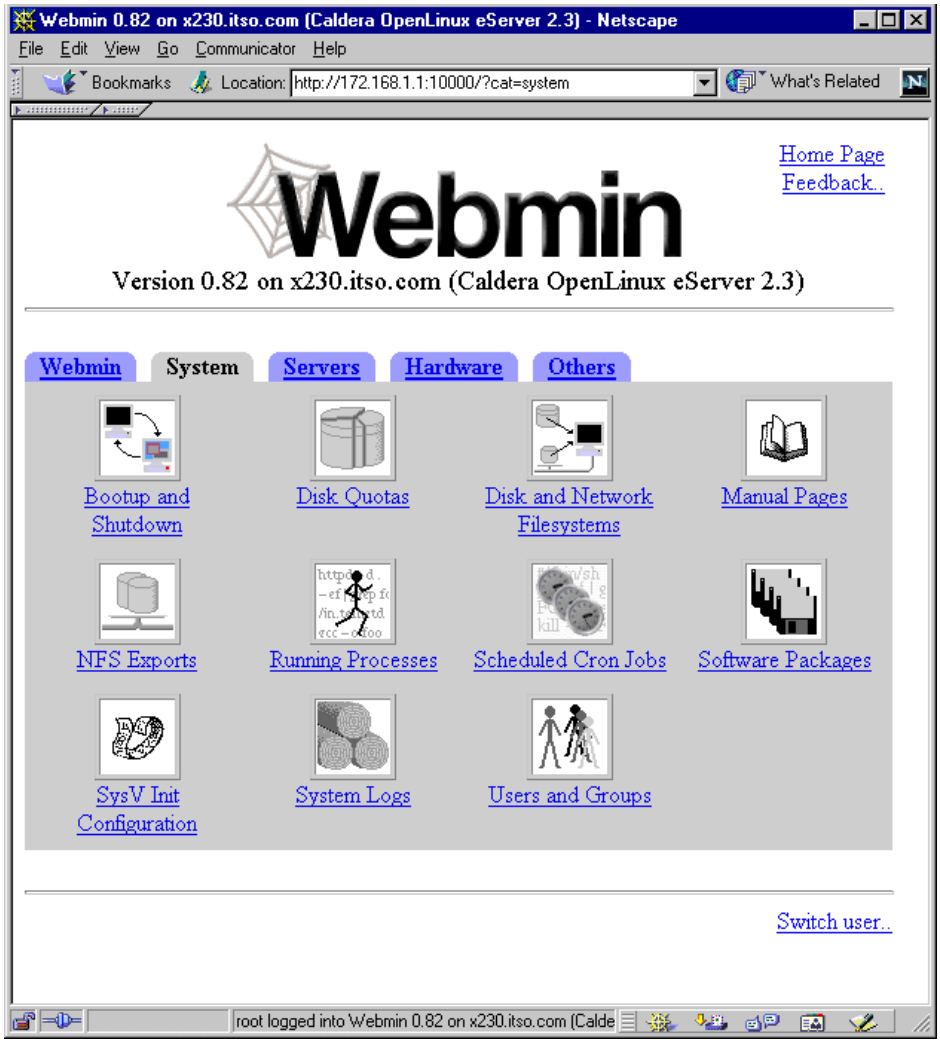

Figure 138. Webmin System tab

On the Webmin System tab, you have the following configuration tasks available:

- Bootup and Shutdown - here you define which services will start at the system startup
- Disk Quotas - if your local filesystem supports quotas, you can configure them here

- Disk and Network Filesystems - here you can mount or unmount the local and networked filesystems
- Manual Pages - here you can search for all Linux manual pages
- NFS Exports - here you can define NFS exports, which are the directories you want to share over the NFS protocol
- Process Manager - here you can manage the running processes
- Scheduled Cron Jobs - here you can create, delete or modify Cron jobs
- Software Packages - here you can manage the installed software packages or install new ones
- SysV Init Configuration - here you can edit the inittab configuration file
- System Logs - here you can see all system log files
- Users and Groups - here you can manage users and groups on your system

3.22.3 Webmin servers

After you click the **Servers** tab in the Webmin main window similar to Figure 137 on page 131, you will see a window similar to Figure 139.

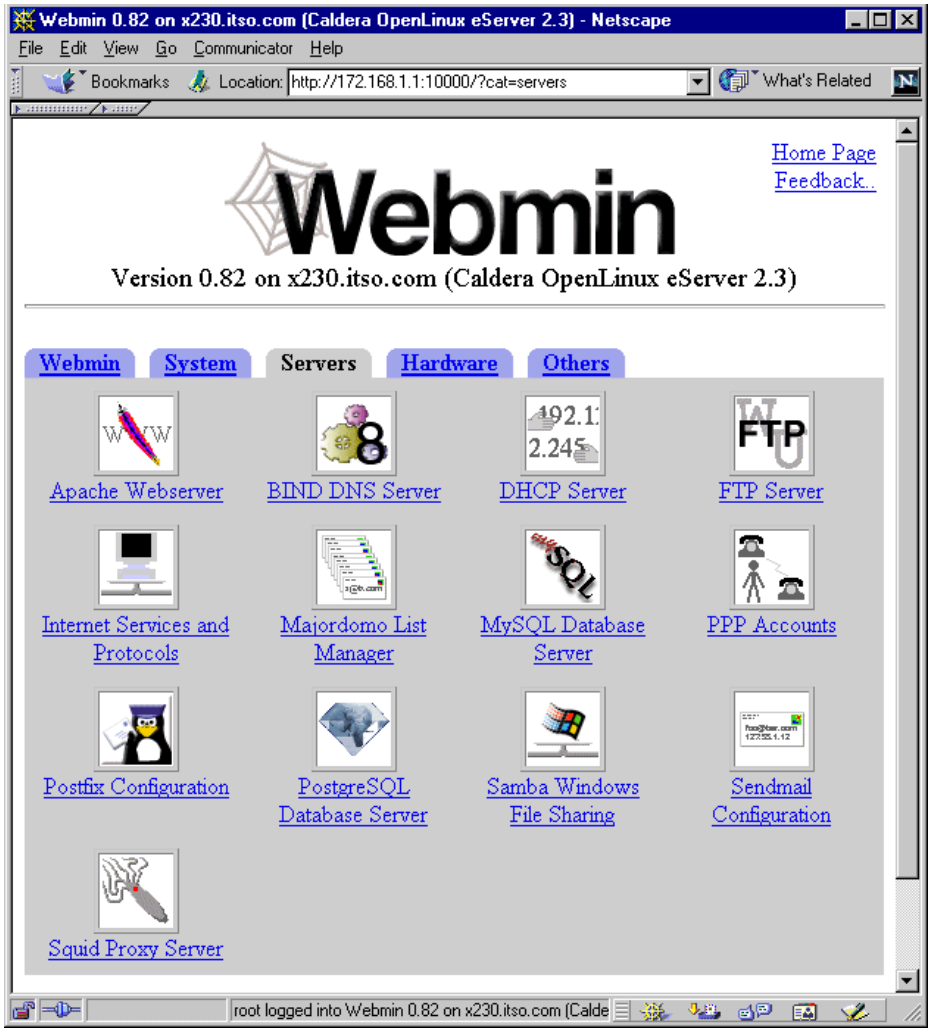

Figure 139. Webmin servers tab

In this section you can configure all the installed server applications on your system.

3.22.4 Webmin hardware

After you click the **Hardware** tab in the Webmin main window similar to Figure 137 on page 131, you will see a window similar to Figure 140.

Figure 140. Webmin hardware tab

On the Webmin Hardware tab, you have the following configuration tasks available:

- Linux Bootup Configuration - here you can configure your LILO bootup configuration
- Linux RAID - if you installed the Linux RAID tools for software RAID, you can configure it here
- Network Configuration - here you can configure your network configuration:
  - Adapters
  - DNS settings
  - Routing and Gateways
  - Host Addresses
- Partition Manager - here you can manage partitions

- Printer Administration - here you can install, configure and remove printers on your Linux server
- System Time - here you can set up the time and date on your server

3.22.5 Webmin Others

After you click the **Others** tab in the Webmin main window similar to Figure 137 on page 131, you will see a window similar to Figure 141.

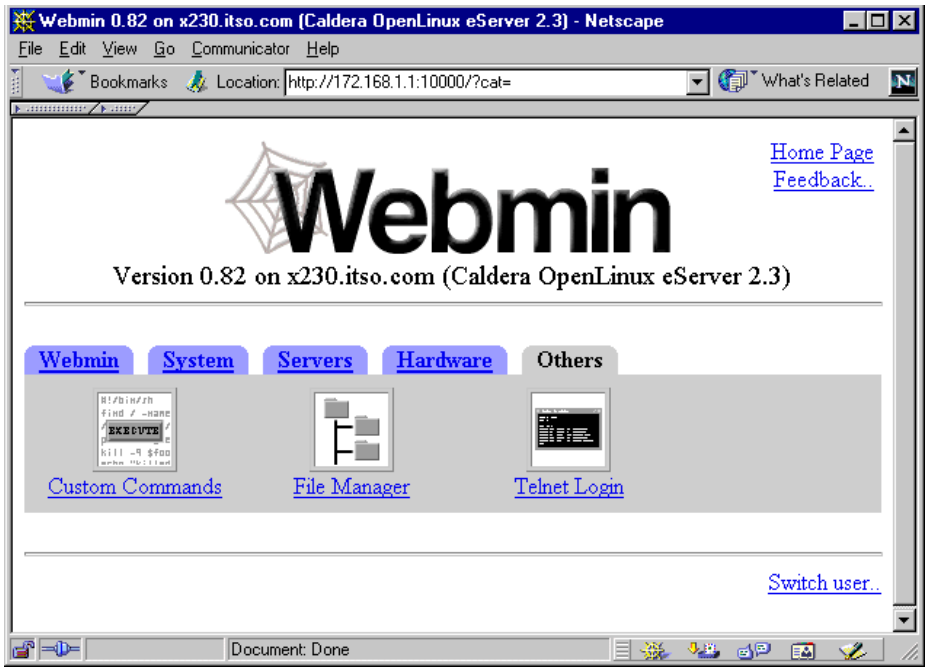

Figure 141. Webmin Others tab

On the Webmin Others tab, you have the following configuration tasks available:

- Custom commands - here you can define custom commands for the Webmin
- File Manager - this is a full-function file manager
- Telnet Login - this a Java Telnet applet allowing you to log in to the system via a Web browser

3.23 Filesystem permissions

Linux has inherent security features, the most noticeable being filesystem 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 by 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 142. Viewing file permissions

Taking the user linux as an example:

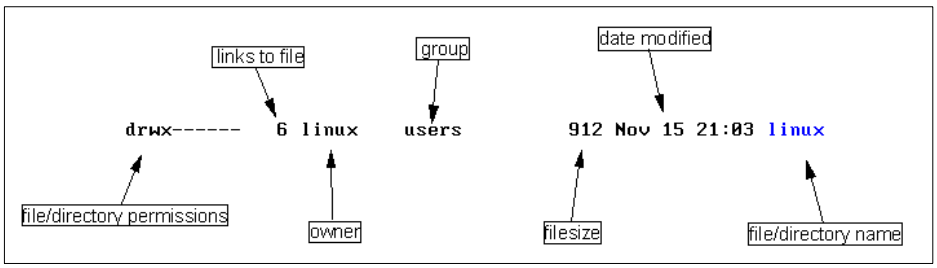

*Figure 143. Explanation of **ls** output*

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 that the file represents a character device.

b - refers to files in the /dev directory. This signifies that 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.

A note about directories: if you set a directory as:

drwxrw-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 into this directory, they will get a `permission denied` error message. This is exactly what happens with regards to users’ home directories.

To change the permissions on a file, use the `chmod` command. Only “root” 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, Here, 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. 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.

4.1 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, since 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 filesystem 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 filesystem 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 filesystem, 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 filesystem 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 filesystem 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.

4.1.1 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.

The Linux /proc filesystem offers a lot of entry points for runtime 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 filesystem, 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 lossy 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`

4.1.2 Kernel recompilation

Recompiling the kernel to only include the drivers and features needed by a machine can help to lower the amount of memory used in the system.

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 compile 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.

4.2 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.

```
 10:31am  up 22 min,  4 users,  load average: 0.00, 0.02, 0.08
69 processes: 67 sleeping, 1 running, 1 zombie, 0 stopped
CPU states:  0.1% user,  1.3% system,  0.0% nice, 98.4% idle
Mem:  127984K av, 109444K used,  18540K free,  81840K shrd,  25888K buff
Swap: 136008K av,      0K used, 136008K free                 35920K cached

  PID USER     PRI  NI  SIZE  RSS SHARE STAT  LIB %CPU %MEM   TIME COMMAND
  864 root      11   0  1100 1100   880 S       0  0.7  0.8   0:00 top
  908 root      11   0  1100 1100   880 R       0  0.5  0.8   0:00 top
  684 root       2   0 10036 9.8M  2032 S       0  0.1  7.8   0:09 X
    1 root       0   0   440  440   376 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
   11 root       0   0     0    0     0 SW      0  0.0  0.0   0:00 scsi_eh_0
  486 root       0   0   588  588   476 S       0  0.0  0.4   0:00 syslogd
  489 root       0   0   816  816   400 S       0  0.0  0.6   0:00 klogd
  500 root       0   0   548  548   460 S       0  0.0  0.4   0:00 inetd
  502 bin        0   0   404  404   328 S       0  0.0  0.3   0:00 rpc.portmap
  522 root       0   0   916  916   752 S       0  0.0  0.7   0:00 amd
  534 root       0   0     0    0     0 SW      0  0.0  0.0   0:00 rpciod
  540 root       0   0   872  872   696 S       0  0.0  0.6   0:00 safe_mysqld
  557 mysql      5   5  1404 1404  1000 S N     0  0.0  1.0   0:00 mysqld
```

Figure 144. Example output of top

`Top` updates the process list in 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 you want to sort the processes by memory consumption, press “m”. To exit from top, press “q”. This will bring you back to the command line.

Similar to `top`, `pstree` displays a hierarchical structure of all currently running processes:

```
[root@x230 /root]# pstree
init-+-amd
     |-atd
     |-cron
     |-6*[getty]
     |-httpd---5*[httpd]
     |-inetd---in.telnetd---login---bash---su---bash---pstree
     |-kdm-+-X
     |     `-kdm---kwm-+-kbgndwm
     |                 |-kfm-+-kvt---bash
     |                 |     |-kvt---bash---top
     |                 |     `-netscape---netscape-statMo---netscape-statMo
     |                 |-kpanel
     |                 |-krootwm
     |                 `-kwmsound
     |-kflushd
     |-klogd
     |-kpiod
     |-kswapd
     |-kupdate
     |-miniserv.pl
     |-named
     |-rpc.mountd
     |-rpc.nfsd---3*[rpc.nfsd]
     |-rpc.portmap
     |-rpc.rstatd
     |-rpciod
     |-safe_mysqld---mysqld---mysqld---mysqld
     |-scsi_eh_0
     |-sendmail
     |-slapd---slapd---slapd
     `-syslogd
```

If you are running a graphical desktop such as KDE, you can also use windows-based tools like KTop, the KDE Task Manager:

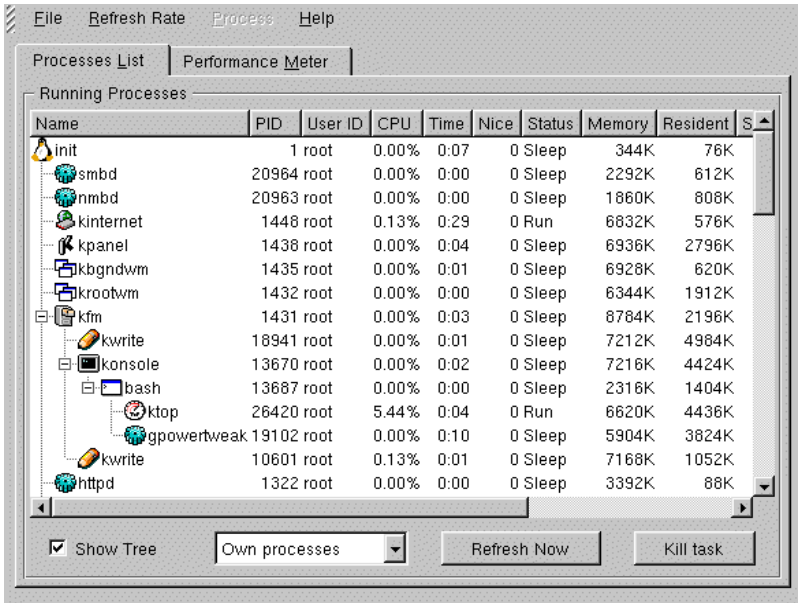

Figure 145. KDE Task Manager: Process List window

KTop offers two different views. It can either display a process list (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.

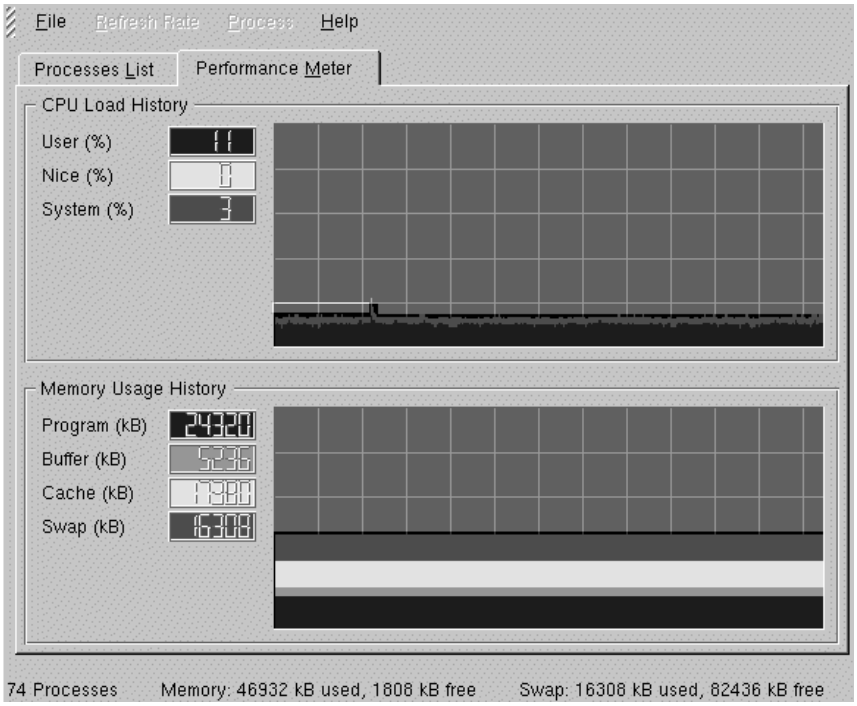

Figure 146. KDE Task Manager: performance meter

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 such as mmx). For example:

```
[root@x230 /root]# cat /proc/cpuinfo
processor       : 0
vendor_id       : GenuineIntel
cpu family      : 6
model           : 7
model name      : Pentium III (Katmai)
stepping        : 3
cpu MHz         : 546.904
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 apic sep mtrr pge mca cmov
pat pse36 mmx fxsr xmm
bogomips        : 1091.17
```

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:

```
[root@x230 /root]# cat /proc/interrupts
           CPU0
  0:     270580          XT-PIC  timer
  1:        923    IO-APIC-edge  keyboard
  2:          0          XT-PIC  cascade
  8:          0    IO-APIC-edge  rtc
 12:      18572    IO-APIC-edge  PS/2 Mouse
 13:          1          XT-PIC  fpu
 14:         14    IO-APIC-edge  ide0
 18:       1990   IO-APIC-level  eth1
 20:      27279   IO-APIC-level  ips
 22:      13236   IO-APIC-level  olympic
 27:        360   IO-APIC-level  PCnet/FAST III 79C975
NMI:          0
ERR:          0
```

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:

```
[root@x230 /root]# cat /proc/ioports
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
01f0-01f7 : ide0
02f8-02ff : serial(auto)
03c0-03df : vga+
03f6-03f6 : ide0
03f8-03ff : serial(auto)
0840-0847 : ide0
0848-084f : ide1
2000-201f : PCnet/FAST III 79C975
2020-203f : Intel Speedo3 Ethernet
4b00-4bff : olympic
```

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:

```
[root@x230 /root]# cat /proc/meminfo
        total:    used:    free:  shared: buffers:  cached:
Mem:  131055616 109469696 21585920 72728576 26509312 36810752
Swap: 139272192        0 139272192
MemTotal:    127984 kB
MemFree:      21080 kB
MemShared:    71024 kB
Buffers:      25888 kB
Cached:       35948 kB
SwapTotal:   136008 kB
SwapFree:    136008 kB
[root@x230 /root]# free
             total       used       free     shared    buffers     cached
Mem:        127984     106912      21072      71088      25888      35960
-/+ buffers/cache:      45064      82920
Swap:       136008          0     136008
```

The /proc/mounts file shows all currently mounted partitions. The `mount` command without parameters will display similar information. For example:

```
[root@x230 /root]# cat /proc/mounts
/dev/root / ext2 rw 0 0
/dev/sda1 /boot ext2 rw 0 0
devpts /dev/pts devpts rw 0 0
/proc /proc proc rw 0 0
x230:(pid522) /auto nfs rw,rsize=1024,wsize=1024,acregmin=0,acregmax=0,acdirmin=
0,acdirmax=0,noac,addr=pid522@x230:/auto 0 0
[root@x230 /root]# mount
/dev/sda5 on / type ext2 (rw)
/dev/sda1 on /boot type ext2 (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/proc on /proc type proc (rw)
x230:(pid522) on /auto type nfs (intr,rw,port=1023,timeo=8,retrans=110,indirect,
map=/etc/am.d/localdev,dev=00000003)
```

The /proc/partitions file displays all existing partitions on all devices. You can also use `fdisk -l` to display this information. For example:

```
[root@x230 /root]# cat /proc/partitions
major minor  #blocks  name

   8     0    8886272 sda
   8     1      23971 sda1
   8     2     136017 sda2
   8     3          1 sda3
   8     5    2240248 sda5
[root@x230 /root]# fdisk -l

Disk /dev/sda: 254 heads, 63 sectors, 1110 cylinders
Units = cylinders of 16002 * 512 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/sda1   *         1         3     23971+  83  Linux
/dev/sda2             4        20    136017   82  Linux swap
/dev/sda3            21      1110   8721090    5  Extended
/dev/sda5            21       300   2240248+  83  Linux
```

The /proc/pci file gives information about all your PCI devices. You can also use the `lspci` command to provide an 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:

```
[root@x230 /root]# cat /proc/pci | more
PCI devices found:
  Bus  0, device   0, function  0:
    Host bridge: Unknown vendor CNB30LE PCI Bridge (rev 4).
      Medium devsel.  Master Capable.  Latency=48.
  Bus  0, device   0, function  1:
    Host bridge: Unknown vendor CNB30LE PCI Bridge (rev 4).
      Medium devsel.  Master Capable.  Latency=48.
  Bus  0, device   1, function  0:
    VGA compatible controller: S3 Inc. Unknown device (rev 1).
      Vendor id=5333. Device id=8904.
      Medium devsel.  Master Capable.  Latency=48.  Min Gnt=4.Max Lat=255.
      Non-prefetchable 32 bit memory at 0xf8000000 [0xf8000000].
  Bus  0, device   2, function  0:
    Ethernet controller: AMD 79C970 (rev 67).
      Medium devsel.  Fast back-to-back capable.  IRQ 27.  Master Capable.  Late
ncy=48.  Min Gnt=24.Max Lat=24.
      I/O at 0x2000 [0x2001].
      Non-prefetchable 32 bit memory at 0xfebffc00 [0xfebffc00].
  Bus  0, device  10, function  0:
    Ethernet controller: Intel 82557 (rev 5).
      Medium devsel.  Fast back-to-back capable.  IRQ 18.  Master Capable.  Late
ncy=48.  Min Gnt=8.Max Lat=56.
      Prefetchable 32 bit memory at 0xfebfe000 [0xfebfe008].
      I/O at 0x2020 [0x2021].
      Non-prefetchable 32 bit memory at 0xfea00000 [0xfea00000].
  Bus  0, device  15, function  0:
    ISA bridge: Unknown vendor Unknown device (rev 79).
      Vendor id=1166. Device id=200.
      Medium devsel.  Master Capable.  No bursts.
  Bus  0, device  15, function  1:
    IDE interface: Unknown vendor Unknown device (rev 0).
      Vendor id=1166. Device id=211.
      Medium devsel.  Master Capable.  Latency=48.
      I/O at 0x840 [0x841].
```

The /proc/swaps file displays information about all active swap partitions. For example:

```
[root@x230 /root]# cat /proc/swaps
Filename                        Type            Size    Used    Priority
/dev/sda2                       partition       136008  0       -1
```

The /proc/version file displays some version information about the Linux kernel. The command `uname -a` will display similar information. For example:

```
[root@x230 /root]# cat /proc/version
Linux version 2.2.17 (root@x230.itso.com) (gcc version egcs-2.91.66 19990314 (eg
cs-1.1.2 release)) #1 SMP Fri Nov 3 15:48:29 EST 2000
[root@x230 /root]# uname -a
Linux x230.itso.com 2.2.17 #1 SMP Fri Nov 3 15:48:29 EST 2000 i686 unknown
```

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 filesystem 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:

```
[root@x220 /]# hdparm -tT /dev/hda

/dev/hda:
 Timing buffer-cache reads:   128 MB in  1.35 seconds =94.81 MB/sec
 Timing buffered disk reads:  64 MB in  3.09 seconds =20.71 MB/sec
[root@x220 /]# hdparm -c1 /dev/hda

/dev/hda:
 setting 32-bit I/O support flag to 1
 I/O support  =  1 (32-bit)
[root@x220 /]# hdparm -tT /dev/hda

/dev/hda:
 Timing buffer-cache reads:   128 MB in  1.30 seconds =98.46 MB/sec
 Timing buffered disk reads:  64 MB in  2.47 seconds =25.91 MB/sec
```

Another popular hard disk performance test is bonnie, found at `http://www.textuality.com/bonnie/`. 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`.

Chapter 5. 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).

5.1 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 - 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).

5.2 Setting up the Samba server

You can check if the Samba package is installed by running kpackage. To start kpackage click the K sign on the panel, select **COAS** and then **kpackage.**

Figure 147. Starting kpackage

When kpackage is started, search for the Server section and then under this find the Network section and expand it. If the Samba package is installed you will see a window similar to Figure 148.

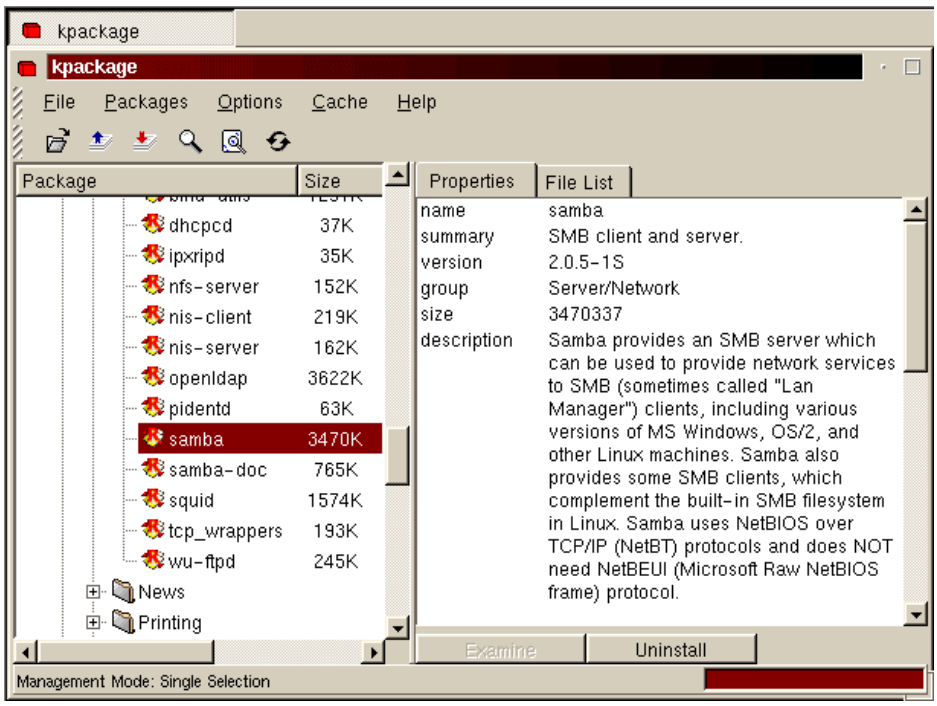

Figure 148. Checking for the Samba package

As you can see in Figure 148, the Samba package is installed.

5.2.1 Configuring the Samba server

In this section we will explain how to configure Samba so it can participate as a file/print server in an existing Window network or be just 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 Caldera OpenLinux the sample configuration file is installed as the /etc/samba.d/smb.conf.sample file.

In Caldera OpenLinux, Samba by default uses the smb.conf file in the directory /etc/samba.d. To begin with, it is enough just to make a copy of the sample file by executing the command:

```
cp /etc/samba.d/smb.conf.sample /etc/samba.d/smb.conf
```

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.
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 define the user's home directories.
   b. Printers - in this subsection you define the available printers.
   c. Shares - this subsection can have more 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 E, "Sample smb.conf SAMBA configuration file" on page 391.

5.2.1.1 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:

```
#=================== Global Settings =======================
[global]
      netbios name = NF5000
      workgroup = LINUX
      server string = Samba Server on Caldera OpenLinux
```

The parameters are described in Table 10.

Table 10. NetBIOS parameters

Parameter	Description
`netbios name`	The Samba server is known by this name on the network. This parameter has the same meaning as the Windows NT computer name. If you do not specify anything it defaults to the server's host name.
`workgroup`	This parameter specifies in which Window NT domain or workgroup the Samba server will participate. It is equivalent to Windows NT domain or workgroup name.

Parameter	Description
server string	This is the description string of the Samba server. It has the same role as the Windows NT description field.

5.2.1.2 Global printing settings

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

```
load printers = yes
printcap name = /etc/printcap
printing = lprng
```

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.
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. Caldera OpenLinux by default uses the LPRNG printing style.

5.2.1.3 Global security settings

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

```
security = user
;   password server = <NT-Server-Name>
encrypt passwords = yes
smb passwd file = /etc/samba.d/smbpasswd
```

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 use the server NetBIOS name.

Parameter	Description
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.
smb passwd file	This parameter tells Samba where encrypted passwords are saved.

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 | /usr/sbin/mksmbpasswd > /etc/samba.d/smbpasswd
```

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

Note

Use the same filename you specified for creating the Samba password file in the smb.conf configuration to tell the Samba server where the password file is.

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 149.

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

Figure 149. Specifying the password for Samba user

Note

Anyone with access to /usr/bin/smbpasswd can change passwords for the Samba users.

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 149. If a Samba user already has defined a password he will need to type the old password before he can change to a new password.

If you want to add a Samba server user later, this 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.

5.2.1.4 Global name resolution settings

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

```
name resolve order = wins lmhosts bcast
wins support = yes
;   wins server = w.x.y.z
```

The parameters are described in Table 13.

Table 13. Name resolution parameters

Parameter	Description
`name resolve order`	With this parameter you specify how the Samba server resolves NetBIOS names into IP addresses. The preferred value is `wins lmhosts bcast`. Refer to the manual page of the smb.conf file for more information.
`wins support`	If this option is enabled the Samba server will also act as a WINS server.
`wins server`	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.

5.2.1.5 Creating shares

In the previous section we explained how to prepare general configuration parameters. But a Samba server can be useful when you offer resources to the 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.

5.2.1.6 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.

5.2.1.7 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 username 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.

5.2.1.8 Creating a printer share

A Samba server uses the same procedure for printer shares as for the home shares. If all share definitions and usernames 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, because 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.

5.2.2 Starting and stopping the Samba server

You can start the Samba server by executing the command:

```
/etc/rc.d/init.d/samba start
```

You will see output similar to the following:

```
[root@nf5000 /root]# /etc/rc.d/init.d/samba start
Starting samba:  smbd nmbd.
```

Figure 150. Samba start window

As you can see, two daemons are started: `smbd` and `nmbd`. `Smbd` is the actual Samba server and `nmbd` is the WINS server.

The Samba server can be stopped by executing the command:

```
/etc/rc.d/init.d/samba stop
```

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

5.2.3 Starting Samba as startup service

You can configure your boot process so Samba is started at the boot. You can do this by using the System Daemon configuration tool. To start the System Daemon configuration tool click the **K** sign on the panel, select **Settings**, then **System** and at the end **Daemons**. You will see a window similar to Figure 151.

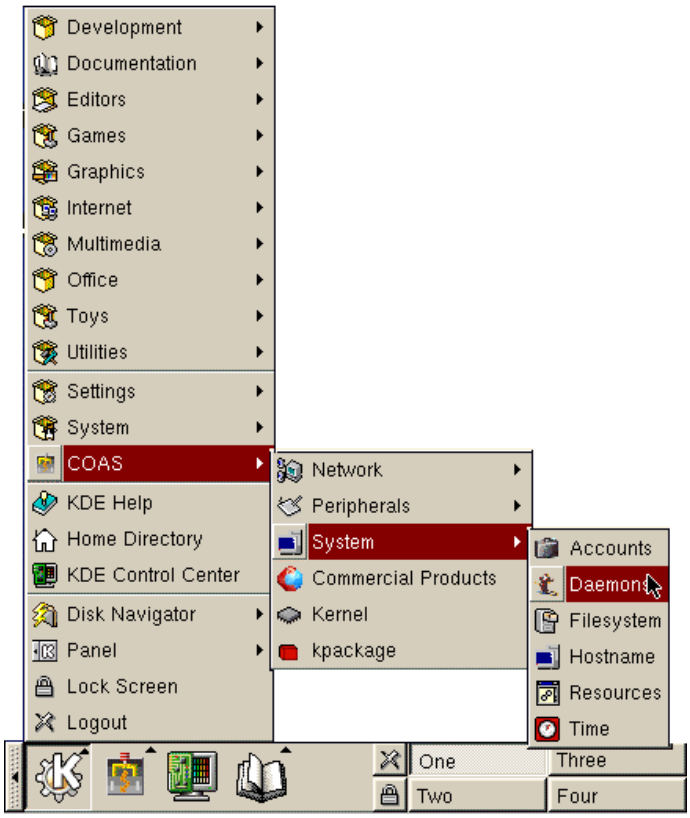

Figure 151. Starting the System Daemon configuration tool

After the System Daemon configuration tool is started you will see a welcome window to COAS administration tools. Click **OK** to continue. You will see a window similar to Figure 152.

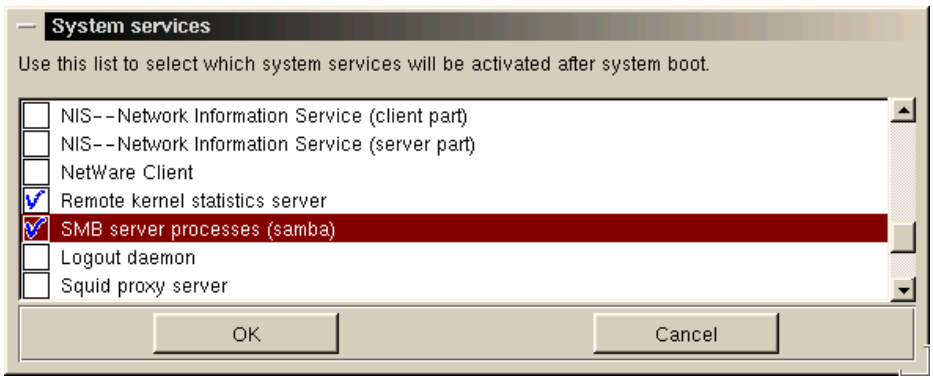

Figure 152. Selecting Samba to start as a boot process

Select **SMB server process (Samba)** on the list. Click **OK** to save your new settings.

When the Linux server is restarted, the Samba server will be started automatically.

5.2.4 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.

Before using SWAT you must check the following:

1. In the /etc/services file you must have the following line:

```
swat 901/tcp
```

2. In the /etc/inetd.conf file you must have the following line:

```
swat stream tcp nowait.400 root /usr/sbin/tcpd swat
```

As you can see, 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
```

If you made any modification to those two files you need to restart inetd. This can be done by executing the commands:

```
/etc/rc.d/init.d/init stop
/etc/rc.d/init.d/init start
```

If you did everything without errors you are ready to use SWAT. 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 153.

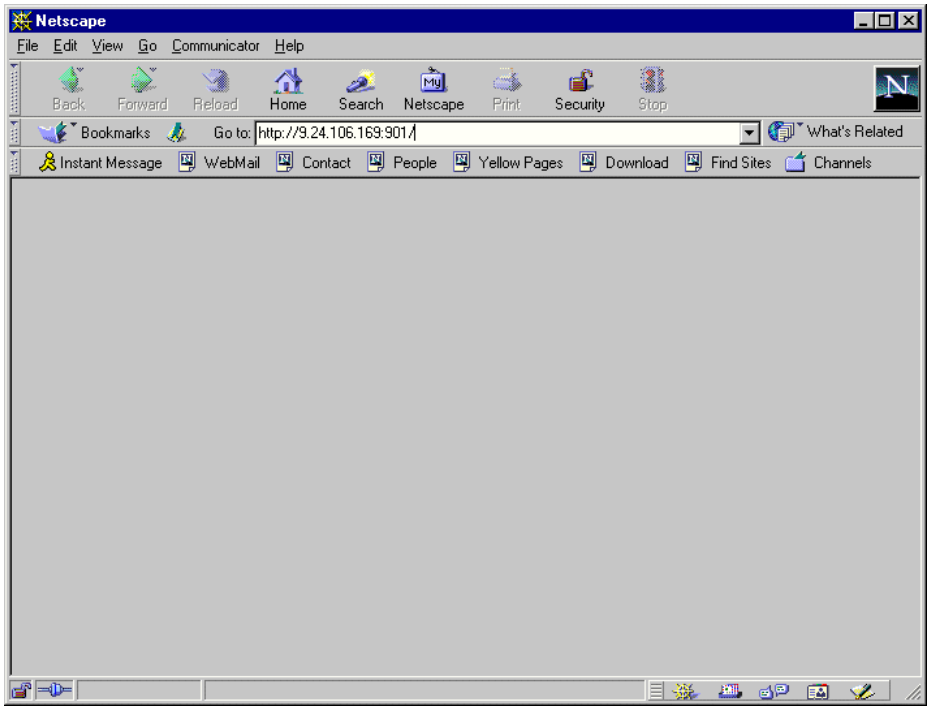

Figure 153. Starting SWAT

After you load the home page of SWAT, you will see a window similar to Figure 154.

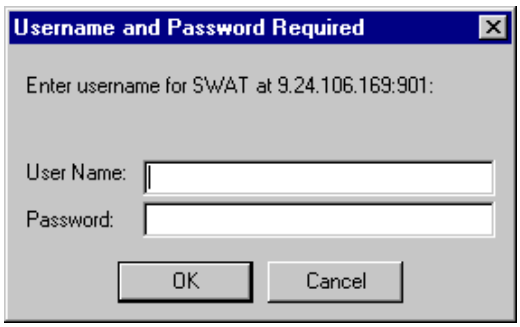

Figure 154. User authorization for SWAT

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 155.

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 do SWAT administration locally only.

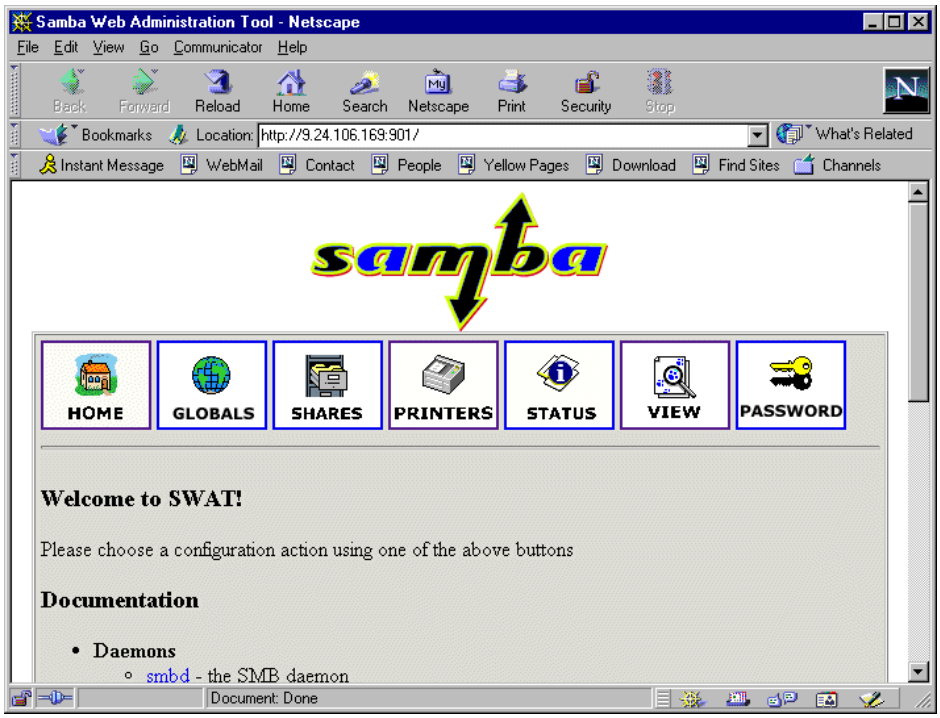

Figure 155. SWAT home page

As you can see in Figure 155, 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.

5.2.4.1 Globals

When you click the **Globals** icon in the main SWAT window, and you will see a window similar to Figure 156.

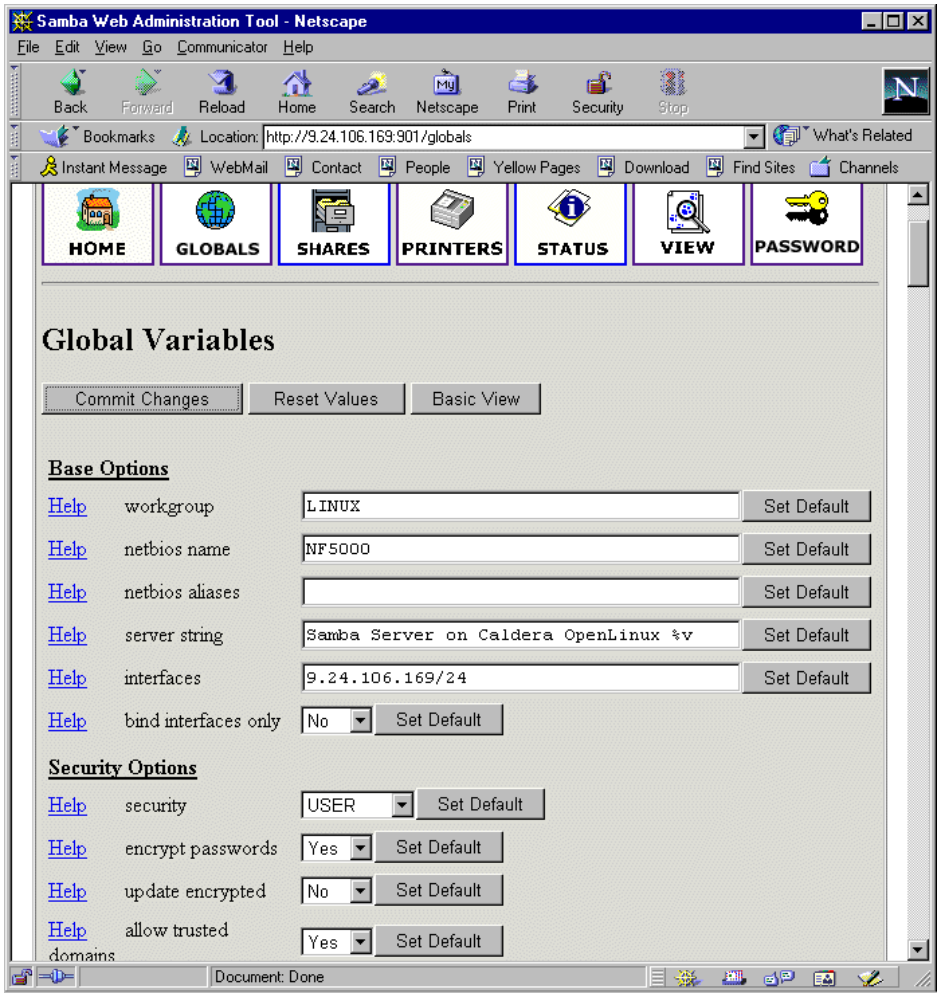

Figure 156. Global section in SWAT

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.** If you get a pop-up window similar to Figure 157, which warns you that you are sending non-secure information over the network, you can easily select **Continue** if you are working locally or if you know that your network is secure.

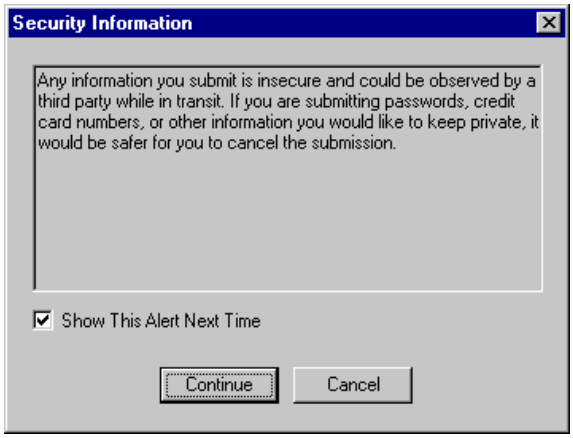

Figure 157. Security warning

5.2.4.2 Shares

When you click the **Shares** icon on any of the SWAT Web pages, you will see a window similar to Figure 158.

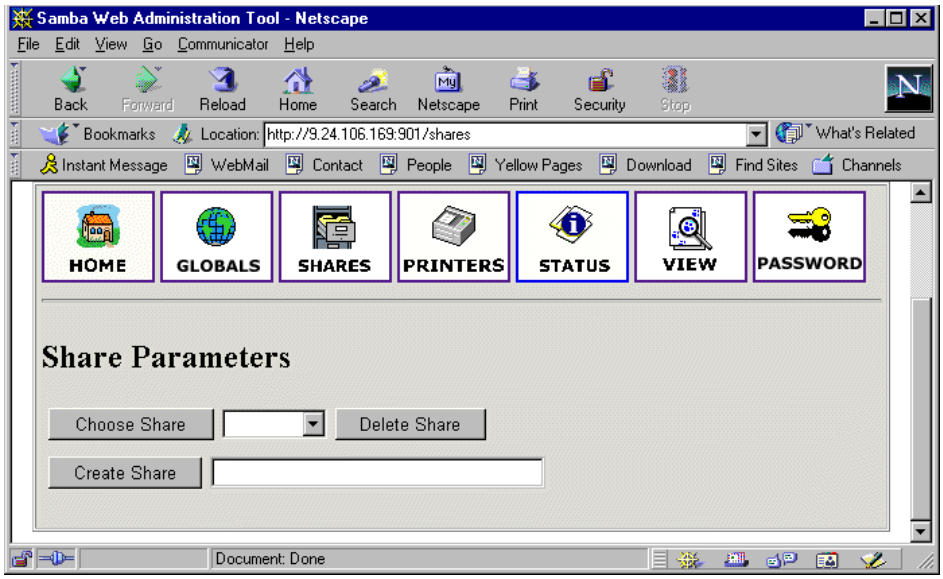

Figure 158. Shares section in SWAT

Here you can:

1. View the defined share

1. Delete share

1. Create a new share

5.2.4.3 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 159.

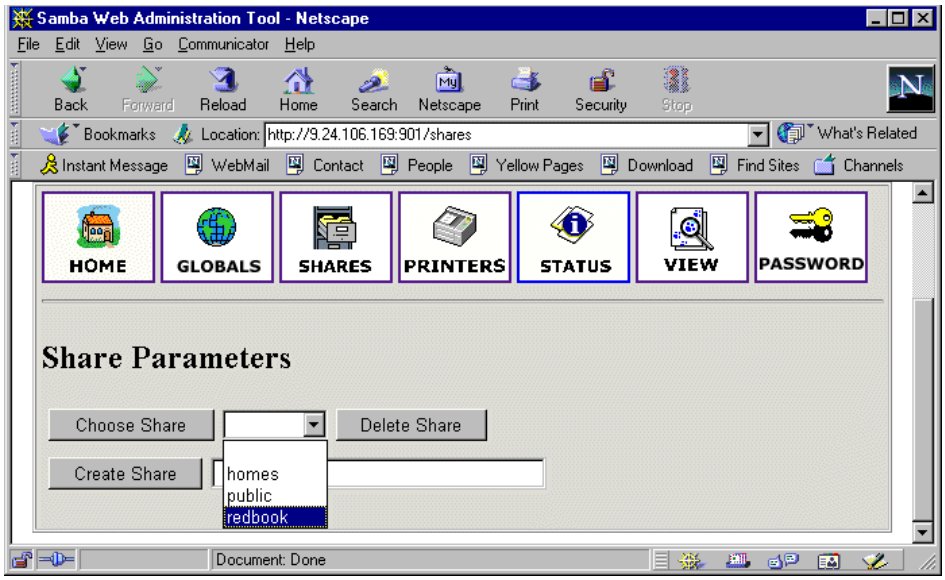

Figure 159. Choosing a share to view

After you have selected the share, click **Choose Share** to view the share properties. You will see a window similar to Figure 160.

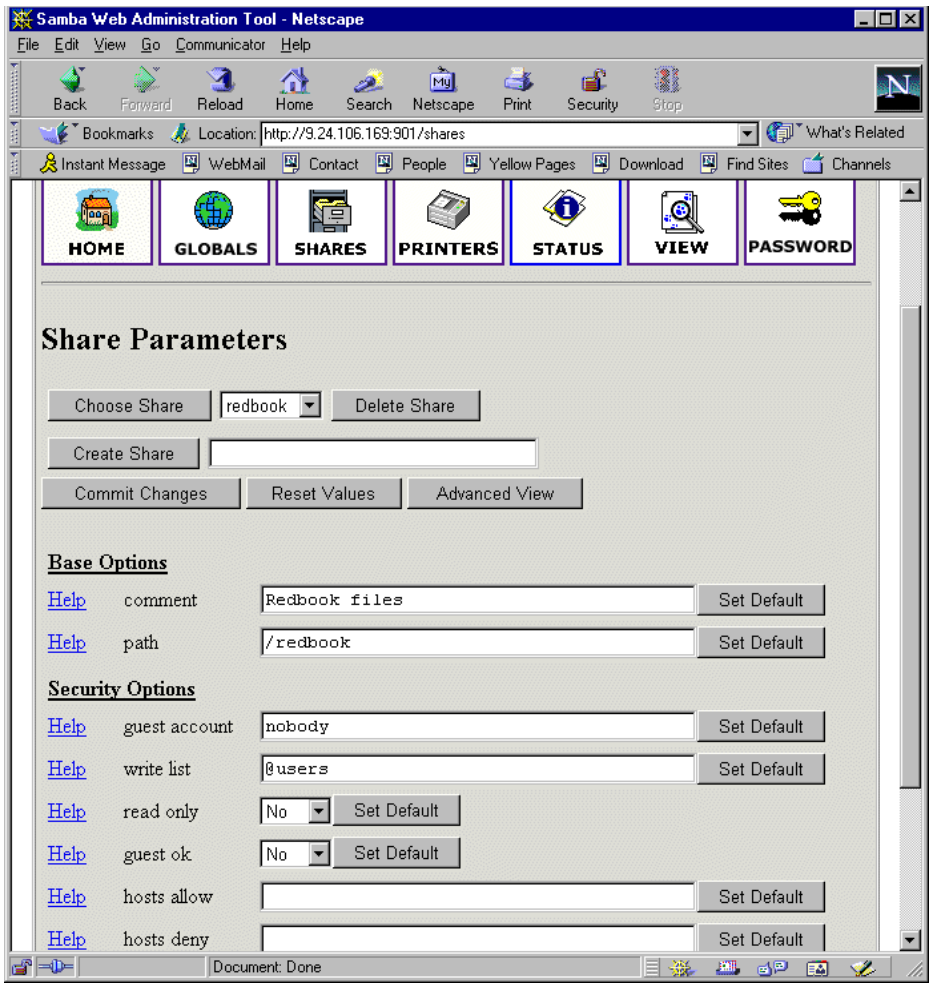

Figure 160. Share properties

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.**

5.2.4.4 Deleting an existing share

To delete an existing share you must first select an already defined share similar to Figure 159. Then click **Delete Share.**

Stop

The share is deleted immediately and without warning.

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

5.2.4.5 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 161.

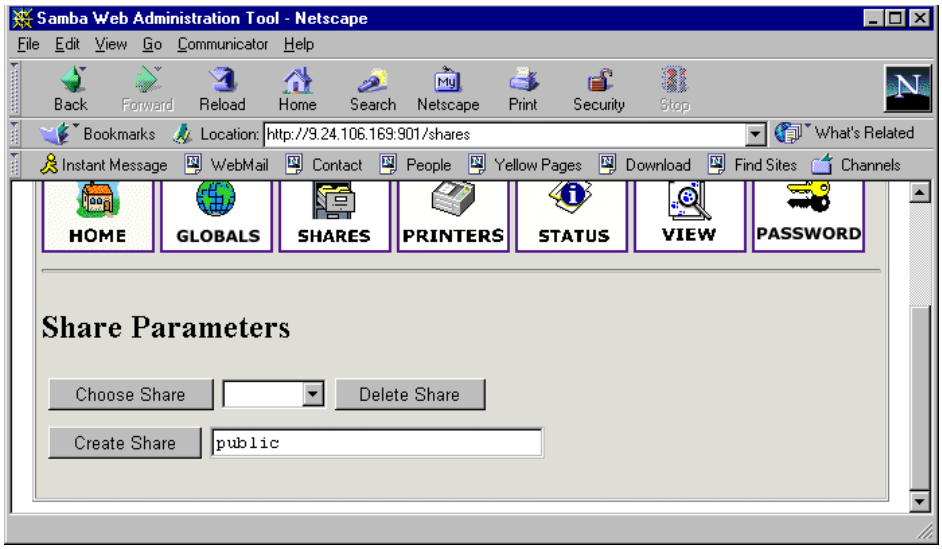

Figure 161. Entering the name for a new share

4. Click **Create Share** to continue, and you will see a window similar to Figure 162.

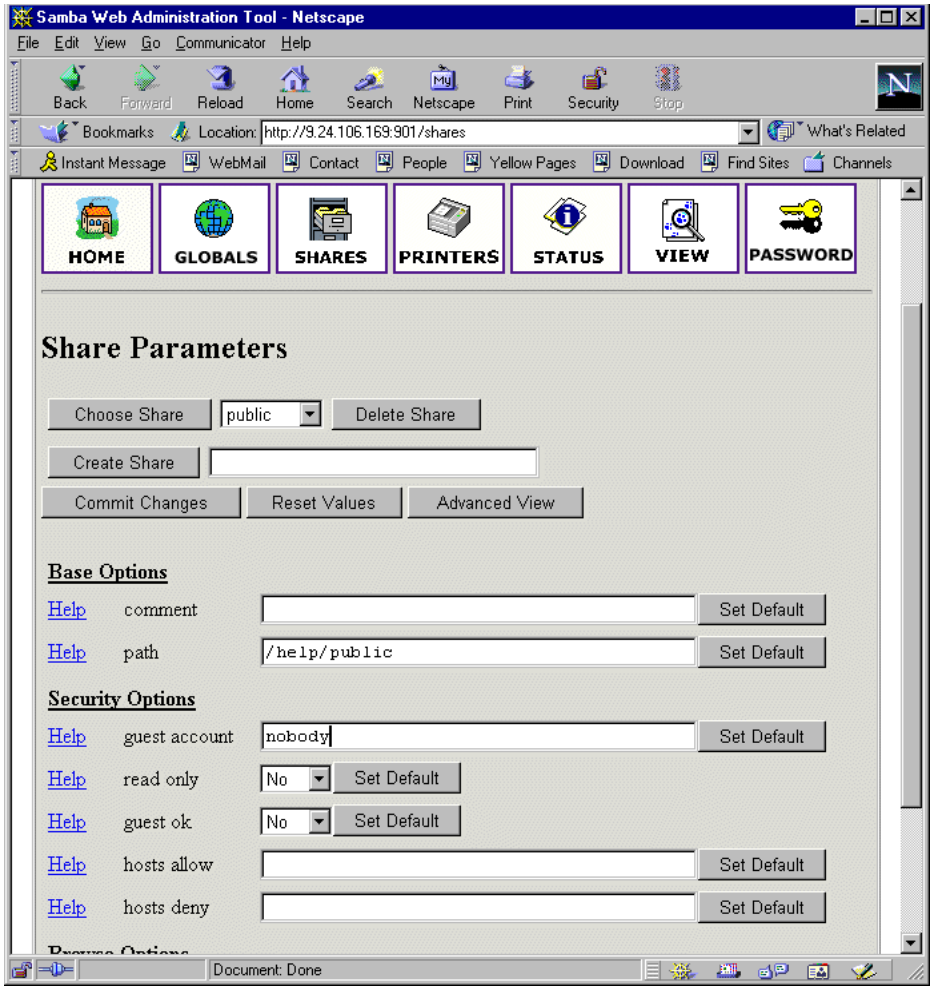

Figure 162. Entering the new share parameters

5. Fill in the needed parameters. If you need to set more advanced parameters, click **Advanced View** and you will see all available parameters. After you typed in all you want, 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 163.

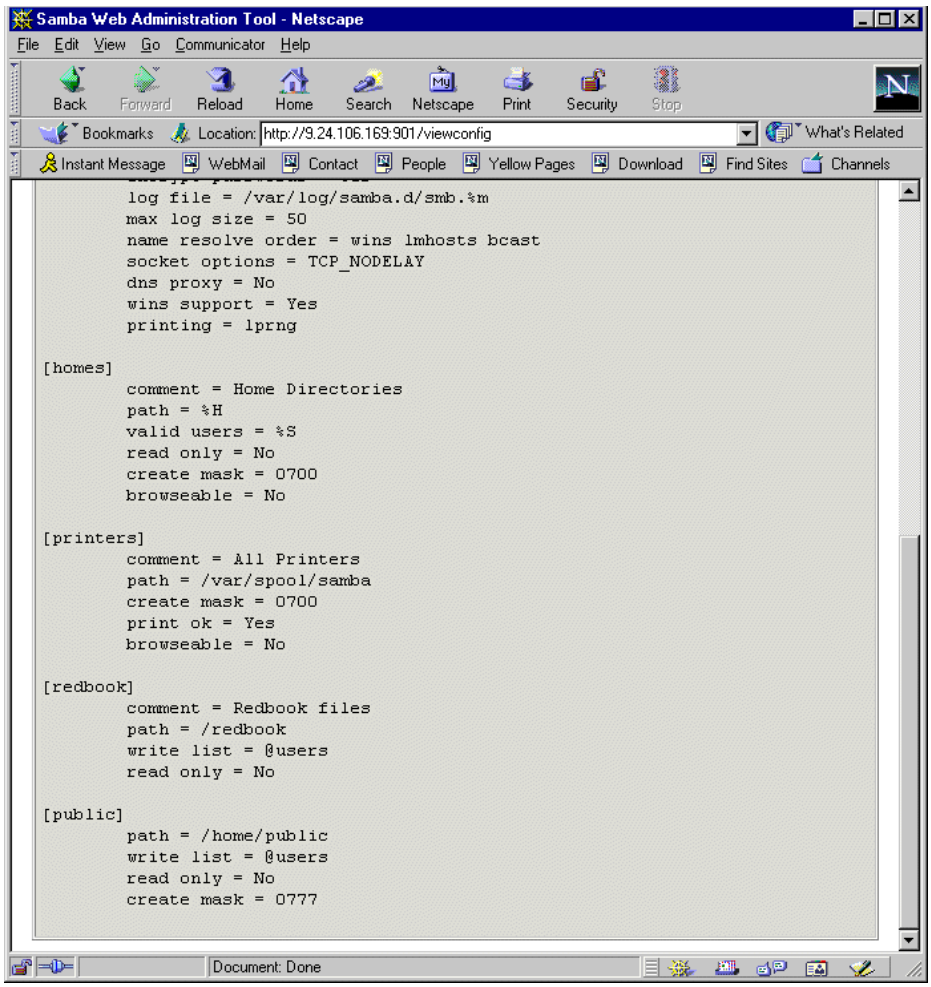

Figure 163. Viewing the smb.conf configuration file

7. Restart the Samba server.

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

5.2.4.6 Restarting the Samba server

The Samba server can be restarted by clicking the **Start** icon on any SWAT Web pages. You will see a window similar to Figure 164.

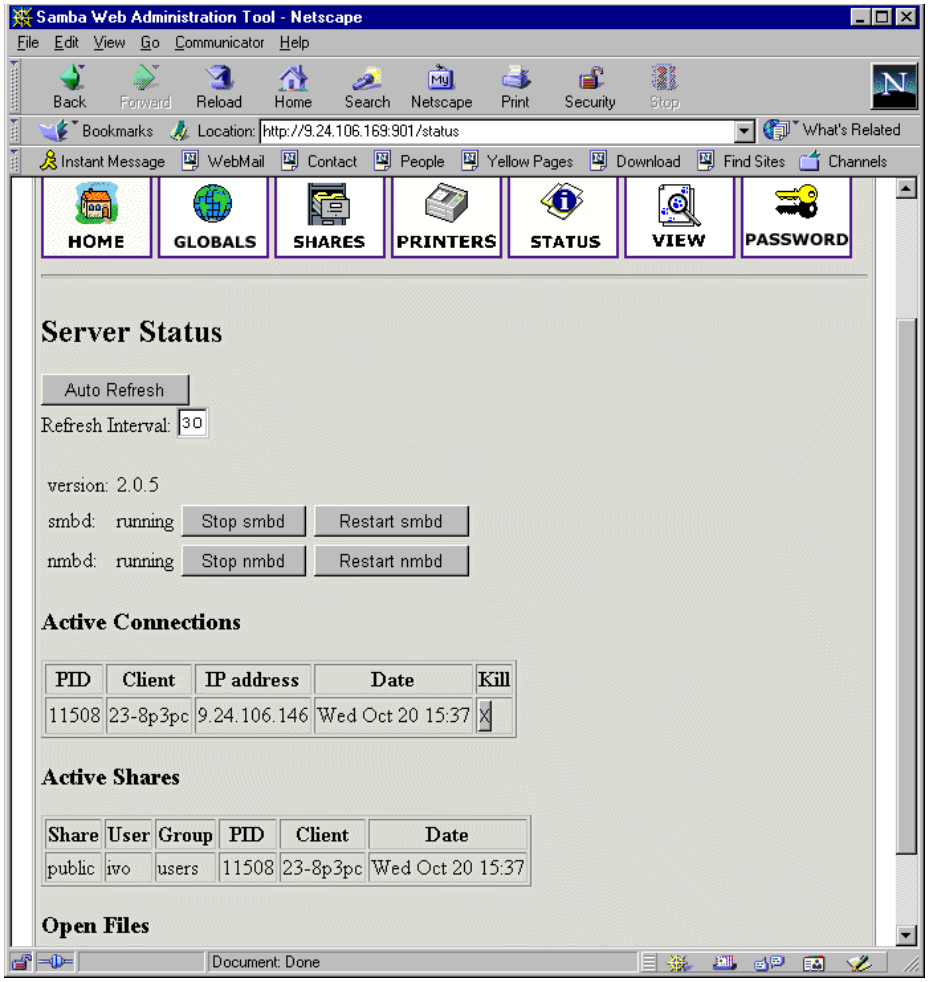

Figure 164. Restarting the Samba server

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

5.2.4.7 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 165.

Figure 165. SWAT printers section

If you want to view the settings for a specific printer, select the printer from the list as you can see in Figure 166.

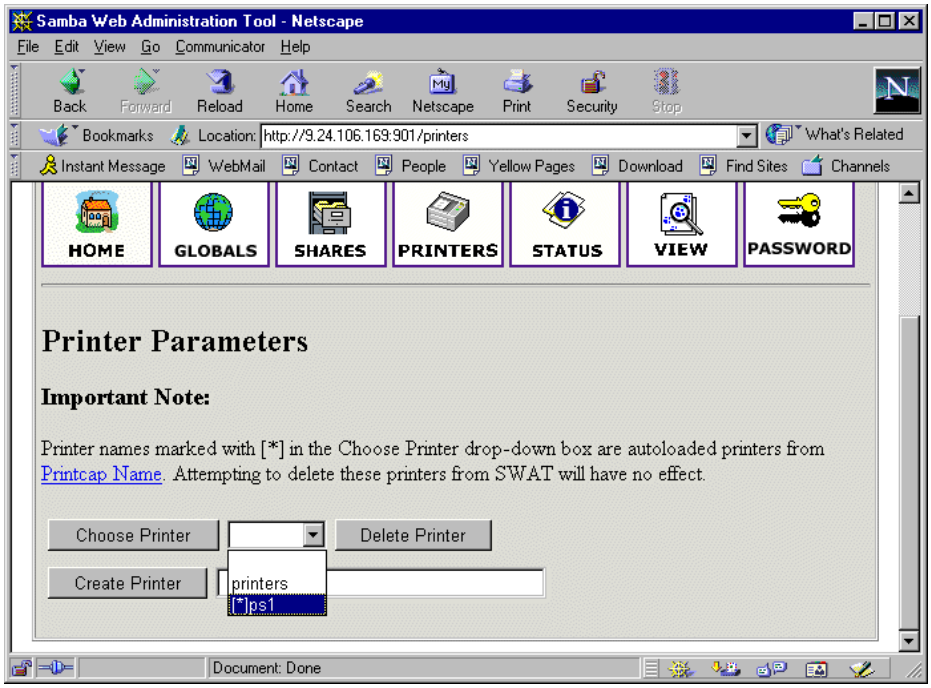

Figure 166. Selecting a printer

After you have selected the printer, click **Choose Printer** to view its properties. You will see a window similar to Figure 167.

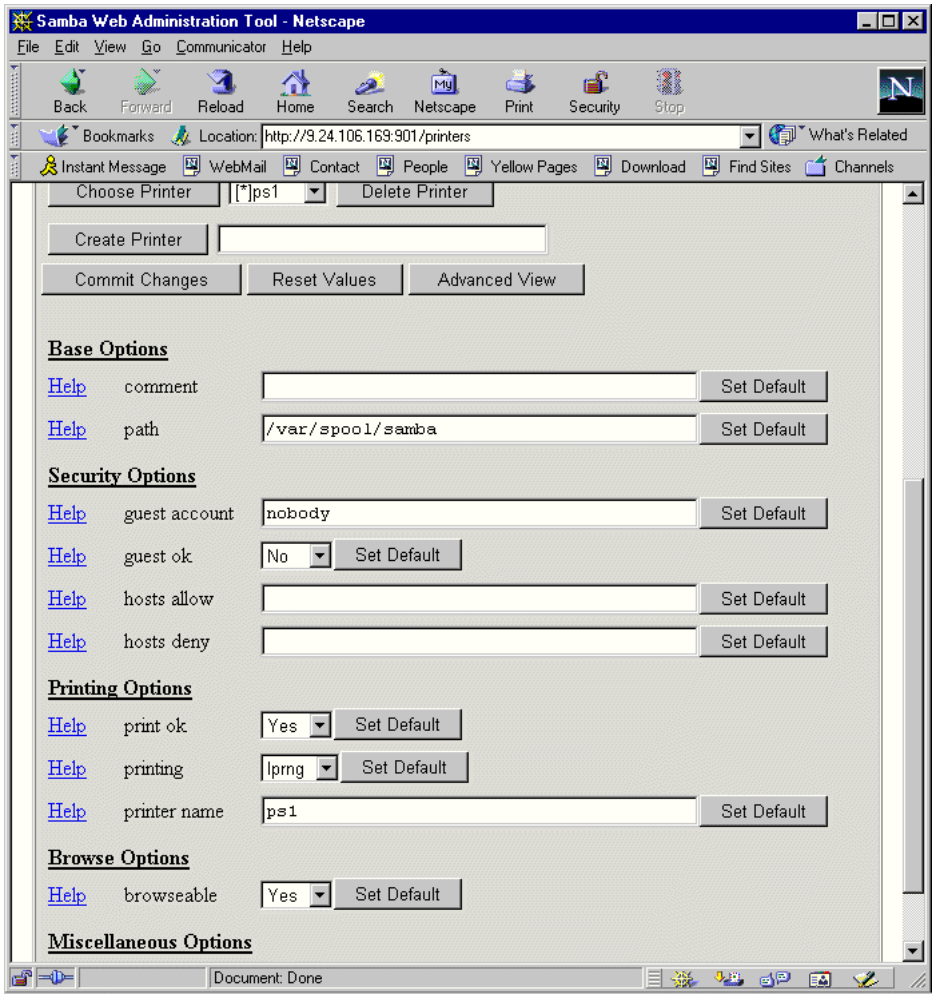

Figure 167. Printer properties

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

5.2.4.8 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. You can access the printer settings by clicking the **Status** icon on the SWAT Web pages, as you can see in Figure 168.

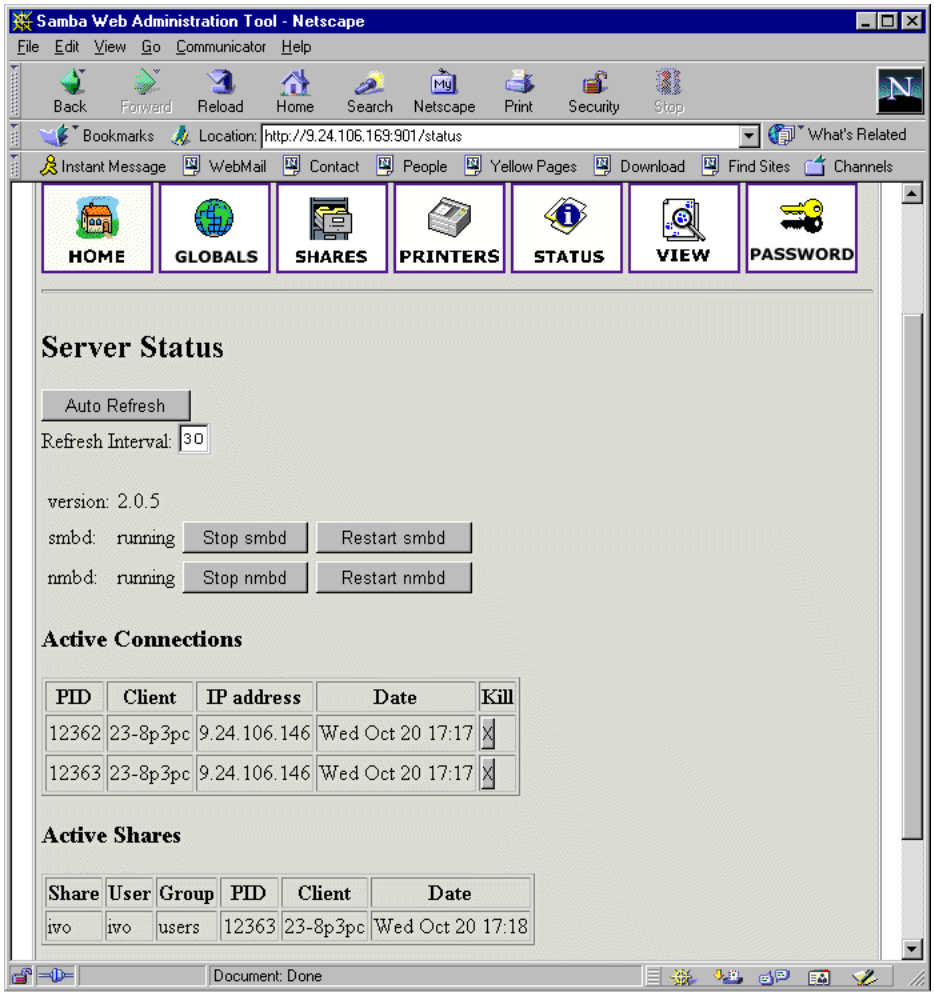

Figure 168. Status section

5.2.4.9 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 169.

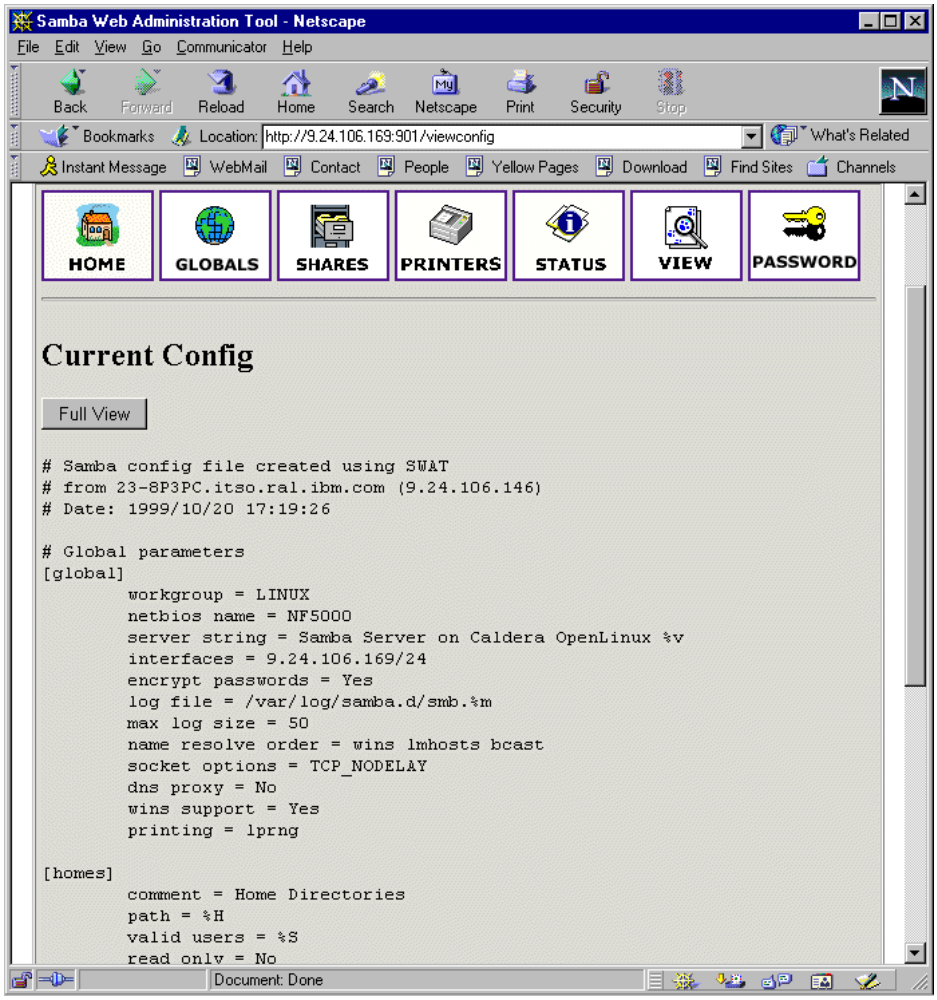

Figure 169. View section of SWAT

5.2.4.10 Password

In this section you can manage the passwords of all Samba users. You can access printer settings by clicking the **Password** icon on the SWAT Web pages similar to Figure 170.

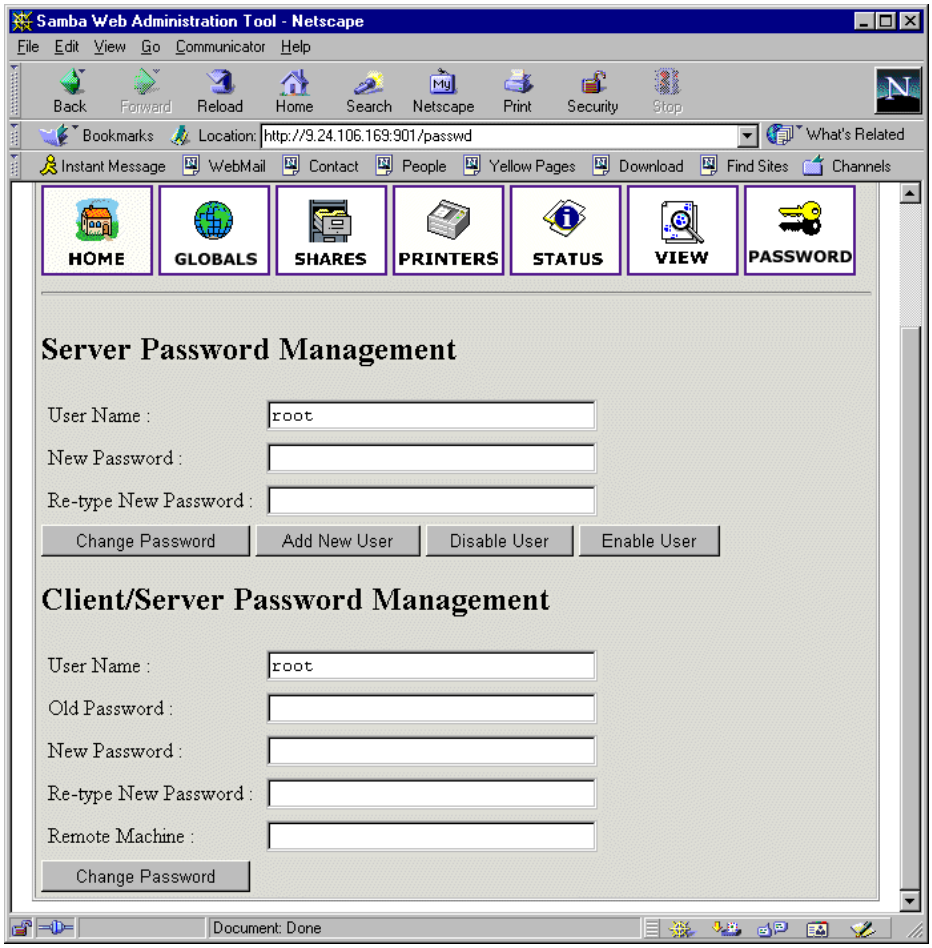

Figure 170. Managing passwords

5.3 Sources of 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 6. 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.

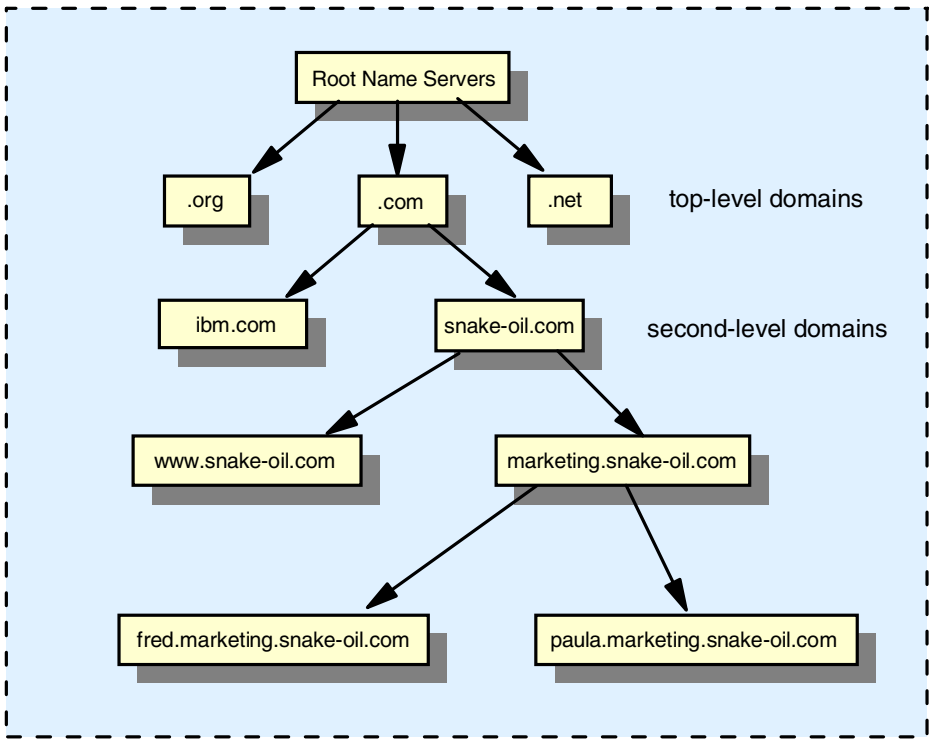

Figure 171. Internet domain hierarchy

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

6.1 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 for bind-8 has been split up into two separate packages in Caldera OpenLinux eServer 2.3: bind, 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 Caldera OpenLinux eServer 2.3, you can use the RPM package manager to query the database of installed packages by entering the following command:

```
rpm -q bind
rpm -q bind-utils
```

If the packages are already installed, RPM will return the version and build number of this packages:

```
bind-8.2.2p4-1
bind-utils-8.2.2p4-1
```

If they are not installed, you will receive the following message:

```
package bind is not installed, or
package bind-utils is not installed.
```

You will then have to install the packages first. Please refer to 3.5, “Adding and removing software packages using kpackage” on page 93 for how to install software packages. These packages are in the Network subsection under the Server section.

6.2 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 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
```

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.

The name server's main configuration file is /etc/named.conf. Open up a text editor and create a new /etc/named.conf according to the following example:

```
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 172. Named.conf 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 173. localhost.zone 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

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

Figure 174. 127.0.0.zone 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:

```
/etc/init.d/named restart
```

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>` or `nslookup`.

If you want DNS to start automatically when the Linux server is started, follow the instructions in 5.2.3, "Starting Samba as startup service" on page 166. But instead of selecting SMB server process (samba), select **Internet domain nameserver (named)**. After you enabled DNS to run as a service it will start automatically when the Linux server boots up.

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:

```
/etc/init.d/named stop
```

Edit the file /etc/named.conf so that it looks like the following example:

```
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 175. Named.conf 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 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 176. Snake-oil.zone 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 signify 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 177. snake-oil.rev file

Now let the name server reload its configuration again by running:

```
/etc/init.d/named 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
Oct 26 18:03:20 ns named[14870]: master zone "snake-oil.com" (IN) loaded (serial 199910
Oct 26 18:03:20 ns named[14870]: master zone "99.168.192.IN-ADDR.APRA" (IN) loaded (ser
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 178. messages file

Your name server should now correctly resolve host names for the snake-oil domain as well.

6.3 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 bind package.

Chapter 7. 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.

7.1 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 though 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.

7.2 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.

7.3 Implementation on Linux

In this section we will discuss how to implement a DHCP server on Linux.

If the DHCP is not yet installed you can install it from the Caldera OpenLinux eServer 2.3 CD.

```
rpm -ivh dhcp2-2.0-0.i386.rpm
```

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.

The DHCP server needs a place to keep track of leases. The /var/state/dhcp/dhcpd.leases file needs to be created to successfully start the DHCP daemon.

```
touch /var/lib/dhcp/dhcpd.leases
```

To start the DHCP daemon, type:

```
/usr/sbin/dhcpd
```

For debugging information, use the -d -f flags.

Chapter 8. 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 the 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).

8.1 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-4.ibm.com/software/webservers/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.

8.2 Apache HTTP Server installation

The Apache HTTP Server is installed and started by default on Caldera OpenLinux eServer 2.3. You can verify that the Apache package is installed by running kpackage as we explained in 3.5, “Adding and removing software packages using kpackage” on page 93. The Apache server package is found by clicking **Server -> WWW**.

The Apache HTTP is usually automatically started on bootup. You can check this by examining the services that will start at server startup, as we described in 3.9, “Daemons (services)” on page 109. If you want Apache HTTP server to start at the server startup, enable the Web server service. To start, stop or reload Apache HTTP server (after a configuration change), run the script:

```
/usr/rc.d/init.d/httpd (start|stop).
```

In Caldera OpenLinux, Apache will serve HTML documents from the directory /home/httpd/html and CGI scripts from /home/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/conf.

If you now point your browser to the server’s IP address, you should see the following start page (/home/http/html/index.html), when the Apache HTTP Server is running:

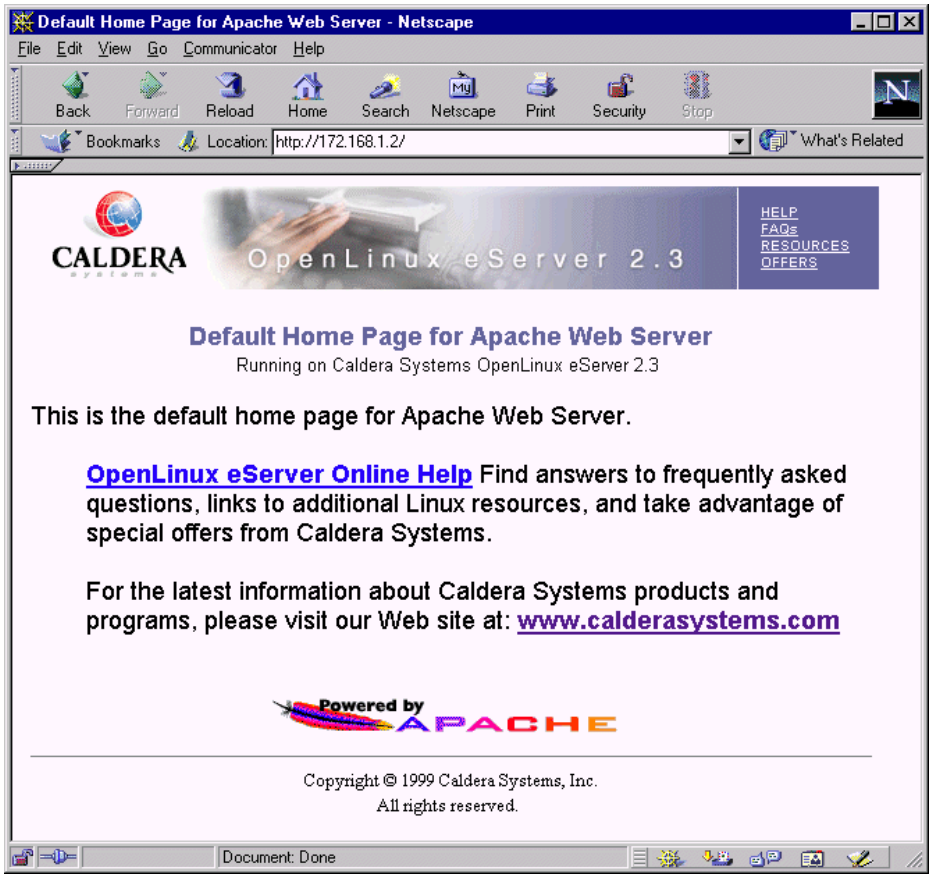

Figure 179. Apache startup page on Caldera OpenLinux eServer 2.3

8.3 IBM HTTP Server installation

To install the IBM HTTP Server on Caldera OpenLinux eServer 2.3, 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-4.ibm.com/software/webservers/httpservers/download.html
```

The HTTPServer.linux.128.tar.gz (or HTTPServer.linux.56.tar.gz for 56-bit encryption) file contains the following packages:

- IBM_HTTP_Server-1.3.12-1.i386.rpm - IBM HTTP Server
- IBM_Apache_Source-1.3.12-1.i386.rpm - Apache 1.3.12-1 source
- IBM_Admin_Server-1.3.12-1.i386.rpm - Administration Server
- IBM_Admin_Server_Forms-1.3.12-1.i386.rpm - Administration Server Web forms
- gsk4bas-4.0-3.57.i386.rpm - Security library
- IBM_SSL_128-1.3.12-1.i386.rpm - 128-bit SSL library

  or

- IBM_SSL_56-1.3.12-1.i386.rpm - 56-bit SSL library
- IBM_SSL_Base-1.3.12-1.i386.rpm - SSL module
- IBM_Machine_Translation-1.3.12-1.i386.rpm - Gateway to IBM MT engine)
- IBM_SNMP-1.3.12-1.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.lnux.128.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:

```
/etc/rc.d/init.d/httpd stop
```

Also, make sure that it will not be started again after the next reboot by disabling the Web server service as we described in 3.9, “Daemons (services)” on page 109.

You now need to install the packages with the following commands (assuming the packages reside in the current directory):

```
rpm -Uvh IBM_HTTP_Server-1.3.12-1.i386.rpm
rpm -Uhv IBM_Admin_Server-1.3.12-1.i386.rpm
```

```
rpm -Uhv IBM_Admin_Server_Forms-1.3.12-1.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:

```
/etc/rc.d/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 window:

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, HTML pages are served from the directory /opt/IBMHTTPD/htdocs and CGI scripts from /opt/IBMHTTPD/cgi-bin. The log files reside in /opt/IBMHTTPD/logs.

8.3.1 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.

If you want this server to be started on bootup, you have to create the correct symbolic links in the directory /sbin/init.d/rc3.d (if you start the system in run level 3, the default run level), or /sbin/init.d/rc5.d (If you use the graphical login, run level 5). You can do this manually with the following commands:

```
cd /etc/rc.d/rc3.d

ln -s ../init.d/ibmhttpd ./S67ibmhttpd
ln -s ../init.d/ibmhttpd ./K01ibmhttpd

cd /etc/rc.d/rc5.d

ln -s ../init.d/ibmhttpd ./S67ibmhttpd
ln -s ../init.d/ibmhttpd ./K01ibmhttpd
```

This will start the IBM HTTP Server in run level 3 and run level 5 and make sure, that it will be properly shut down when switching into another run level (for example shutdown).

8.3.2 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:

```
[root@x220 /root]# /opt/IBMHTTPServer/bin/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] -> wwwgroup

************************************************************
Please supply the Directory containing the files for
which a change in the permissions is necessary.
************************************************************
[default: /opt/IBMHTTPServer/conf] ->

These are the file(s) and directory for which we will be changing
Group permissions:

-rw-r--r--   1 root     root         4359 Sep  6 08:37 admin.conf
-rw-r--r--   1 root     root         4359 Sep  6 08:37 admin.conf.default
-rw-r--r--   1 root     root         7453 Sep  6 08:37 admin.msg
-rw-r--r--   1 root     root            1 Sep  6 08:37 admin.passwd
-rw-r--r--   1 root     root        31144 Sep  6 08:15 httpd.conf
-rw-r--r--   1 root     root        31144 Sep  6 08:15 httpd.conf.default
-rw-r--r--   1 root     root        48900 Sep  6 08:15 httpd.conf.sample
-rw-r--r--   1 root     root        12441 Sep  6 08:15 magic
```

```
-rw-r--r--   1 root     root        12441 Sep  6 08:15 magic.default
-rw-r--r--   1 root     root         9957 Sep  6 08:15 mime.types
-rw-r--r--   1 root     root         9957 Sep  6 08:15 mime.types.default
drwxr-xr-x   2 root     root         1024 Nov  4 13:38 /opt/IBMHTTPServer/conf

************************************************
CONTINUE - Perform Changes ENTER 1
QUIT -     No Changes      ENTER 2
************************************************
[default: QUIT - 2] -> 1

>>>Validating Group Name: 'nogroup'<<<
        Group Name: 'nogroup' already exists

>>>Validating UserID:wwwrun<<<
        UserID: 'wwwrun' already exists

>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
Changing Group:
CMD: 'chgrp  nogroup /opt/IBMHTTPServer/conf /opt/IBMHTTPServer/conf/* '
<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<


>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
Changing Group permissions:
CMD: 'chmod  g+rw /opt/IBMHTTPServer/conf /opt/IBMHTTPServer/conf/*'
<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<

Here are the new file(s) and directory permissions:

drwxrwxr-x   2 root     wwwgroup     1024 Nov  4 13:38 /opt/IBMHTTPServer/conf
-rw-rw-r--   1 root     wwwgroup     4359 Sep  6 08:37 /opt/IBMHTTPServer/conf/a
dmin.conf
-rw-rw-r--   1 root     wwwgroup     4359 Sep  6 08:37 /opt/IBMHTTPServer/conf/a
dmin.conf.default
-rw-rw-r--   1 root     wwwgroup     7453 Sep  6 08:37 /opt/IBMHTTPServer/conf/a
dmin.msg
-rw-rw-r--   1 root     wwwgroup        1 Sep  6 08:37 /opt/IBMHTTPServer/conf/a
dmin.passwd
-rw-rw-r--   1 root     wwwgroup    31144 Sep  6 08:15 /opt/IBMHTTPServer/conf/h
ttpd.conf
-rw-rw-r--   1 root     wwwgroup    31144 Sep  6 08:15 /opt/IBMHTTPServer/conf/h
ttpd.conf.default
-rw-rw-r--   1 root     wwwgroup    48900 Sep  6 08:15 /opt/IBMHTTPServer/conf/h
ttpd.conf.sample
-rw-rw-r--   1 root     wwwgroup    12441 Sep  6 08:15 /opt/IBMHTTPServer/conf/m
agic
-rw-rw-r--   1 root     wwwgroup    12441 Sep  6 08:15 /opt/IBMHTTPServer/conf/m
agic.default
-rw-rw-r--   1 root     wwwgroup     9957 Sep  6 08:15 /opt/IBMHTTPServer/conf/m
ime.types
-rw-rw-r--   1 root     wwwgroup     9957 Sep  6 08:15 /opt/IBMHTTPServer/conf/m
ime.types.default
Changes Completed

****************************************************************
Configuration file: '/opt/IBMHTTPServer/conf/admin.conf'
```

```
will be saved as '/opt/IBMHTTPServer/conf/admin.conf.19:02:53_309'
Do you wish to update the Administration Server Configuration file
CONTINUE  enter 1
EXIT      enter 2
*****************************************************************
[default: QUIT - 2] -> 1
USER DONE
GRoup  DONE
Successfully updated configuration file
Old configuration file saved as '/opt/IBMHTTPServer/conf/admin.conf.19:02:53_309
[root@x220 /root]#
```

To summarize the above steps: the Administration Server will be running under the user name “wwwrun” and the group “wwwgroup”.

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 8.3.1, “Activating IBM HTTPD on system bootup” on page 205, 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 204, you need to enter the user name and password you defined for the Administration Server user. If entered correctly, you will see the welcome window of the Administration Server:

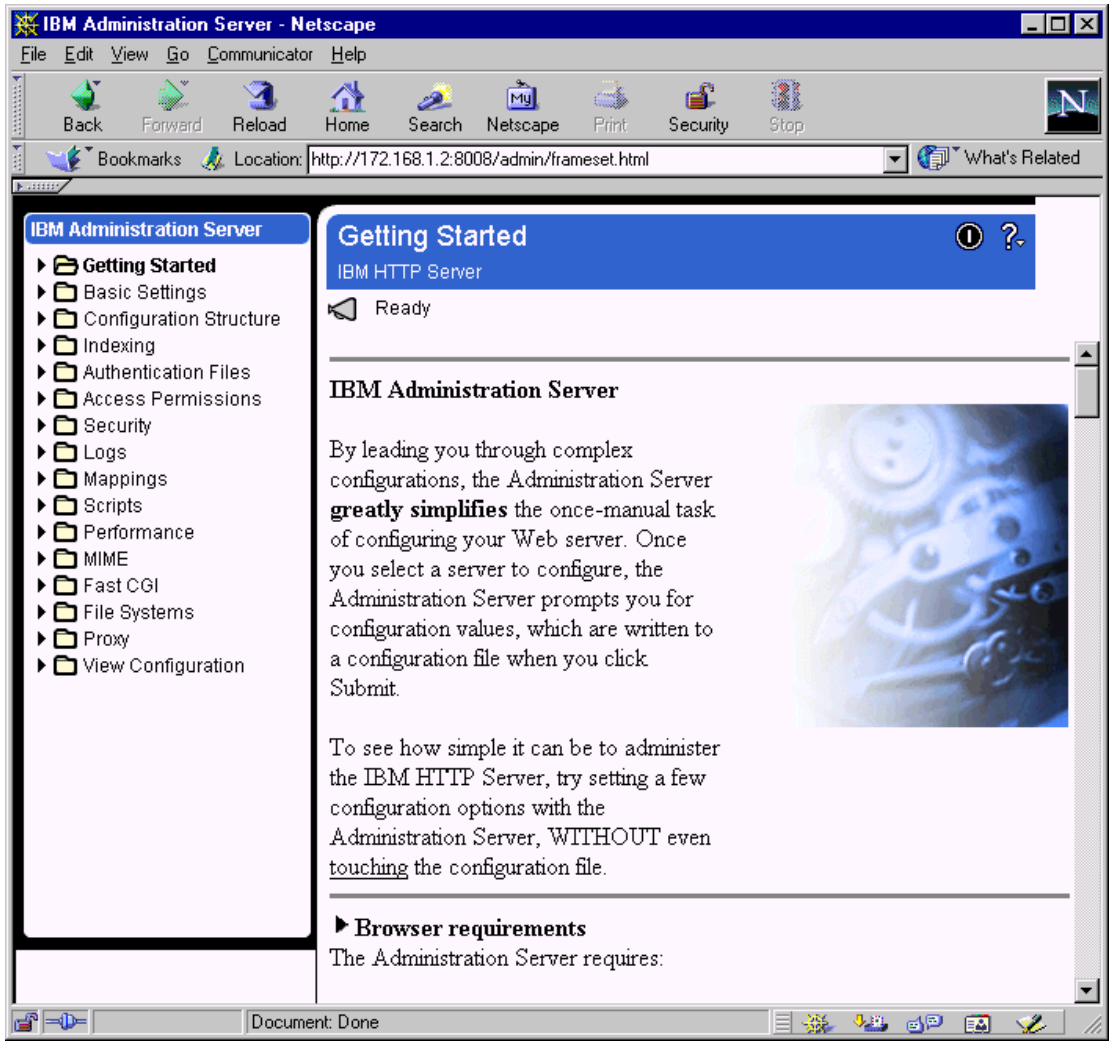

Figure 181. Administration Server startup window

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 about the different configuration options.

8.4 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 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 the Max Clients value 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 containg 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)
```

Chapter 9. sendmail

Communicating with other people is one of the most desirable experiences in life. Sending electronic mail is a way to communicate with people all over the globe. Electronic mail can be more reliable, cheaper and faster than ordinary mail.

9.1 What is sendmail?

As you can tell from the name, sendmail is used to send mail. However, sendmail is not sending old-fashioned mail, but electronic mail, which becomes more important every day. But in spite of that, sendmail is basically acting as a post office. It receives mail from a sender and passes the mail on to the recipient post office. At the recipient post office a local postman delivers the mail to the recipient mailbox. Sendmail is a powerful Mail Transport Agent (MTA) and is used to pass the mail to another MTA, which can be sendmail or some other application capable of handling electronic mail.

9.2 What you can do with sendmail

With sendmail your Linux server can become a server for electronic mail. You can handle mail for users of a Linux server locally and users do not have to ask for mail accounts. The users on your Linux server will have their mailboxes locally and they will still be able to send mail to people anywhere. When you set up sendmail, you can also offer mail service to the users which have accounts on the other network servers, which do not provide Internet mail service.

9.3 Before you begin

In the following sections we explain how to set up a mail server on your Linux server. In this explanation we will use our lab network setup, and all setup is related to this lab setup. You can easily adapt this to your existing installation. You can see our lab network in Figure 182. Before you proceed with the implementation, you need to verify if these packages are installed:

- bind-8.2.2p4-1
- bind-utils-8.2.2p4-1
- sendmail-8.9.3-1
- sendmail-cf-8.9.3-1

You can do this by starting the kpackage tool from COAS tools as we explained in 3.5, “Adding and removing software packages using kpackage” on page 93. These packages are in the Network and Mail subsections under the Server section. If the packages are not installed, refer to the same package for instructions on how to install them.

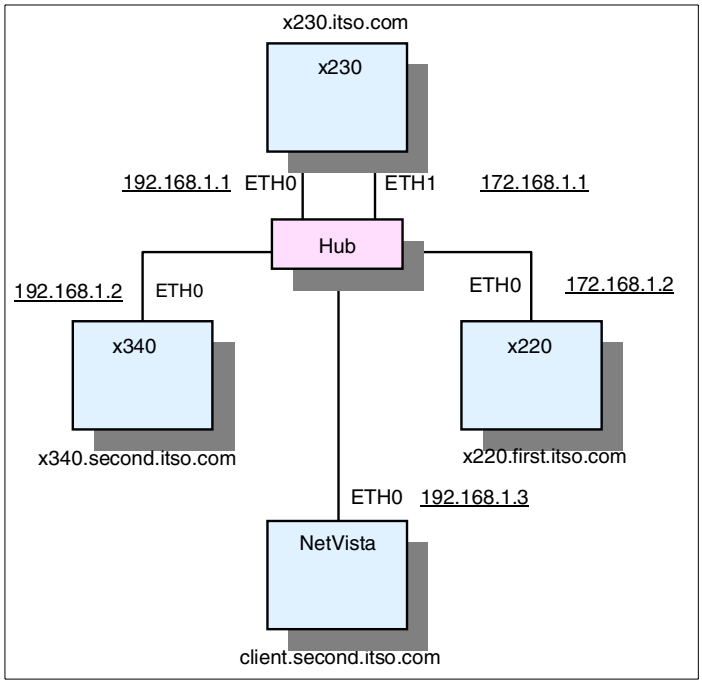

Figure 182. Lab network installation for sendmail setup

Note

For successful operation of sendmail you need the correct settings for the Domain Name System (DNS) server. This means that you have to set up your own DNS correctly or have access to another DNS.

As you can see in Figure 182 our network consists of three domains:

1. Itso.com - this is the master domain. All computers in this domain have a .itso.com extension. In this domain we have a server running the master DNS.

2. First.itso.com - this is the first subdomain of the .itso.com domain. The main server in this domain is running the DNS for this domain and is also the mail (sendmail) server for the users of the domain.
3. Second.itso.com - this is the second subdomain of the .itso.com domain. The main server in this domain is running the DNS for this domain and is the mail (sendmail) server for the users of the domain.

All users of the mail server have user name/password definition on their domain server even if they are using other physical servers or workstations. They need this user name/password for accessing the mail. Each defined user has a mailbox on the mail server. He can reach his mailbox over the network with his client connecting to the mail server. When the connection is established, the user can download his messages to his workstation and delete them from the server. Or he can remotely connect to the server and read his mail on the server, but in this case the mail stays in the mailbox on the server.

When the users are using the server for mail only and they have their own workstations, the mail servers are set up with limited space for each mail box. That means that users have to download their mail regularly from the server to make room for the new messages. In the environment where users use the server for their operations the mailboxes are usually bigger. Users can still reach their mail remotely with the client. In that case they do not download the messages.

The most commonly used protocols for sending and delivering mail are SMTP and POP3. The Simple Mail Transfer Protocol (SMTP) is used for sending mail from the mail client and the Post Office Protocol (POP3) is for getting the mail from the mail server. Along with all other protocolsm sendmail also supports the SMTP and POP3 protocols. In our setup we used SMTP/POP3 protocols for the mail exchange.

Before we start describing how to set up a mail server, we will describe how to set up DNS for our lab network. That is because the correct DNS setup is important for successful operation of the mail server.

9.4 Network configuration

Each subdomain is on its own network. The “first.itso.com” domain is on the network 172.168.1.0 and “second.itso.com” domain is on the network 192.168.1.0. The server running the “itso.com” domain is acting as a gateway for both subdomains.

We have the following network definitions:

1. The gateway server with two Network Interface Cards (NIC). The first NIC has an IP 192.168.1.1. The second NIC has the IP 172.168.1.1. The server has enabled IP Forwarding so it can act as a gateway for subnetworks. This server is also the DNS server for the “itso.com” domain.
2. The server for the “first.itso.com” domain has one NIC with IP address 172.168.1.2.
3. The server for the “second.itso.com” domain has one NIC with IP address 192.168.1.2.
4. The client in the “second.itso.com” domain has one NIC with IP address 192.168.1.3.
5. The client in the “first.itso.com” domain has one NIC with IP address 172.168.1.3.

9.5 Setting up the DNS configuration

As mentioned before our lab network setup was created with a master domain called “.itso.com” with two subdomains called “.first.itso.com” and “.second.itso.com”. You can see this setup in Figure 183.

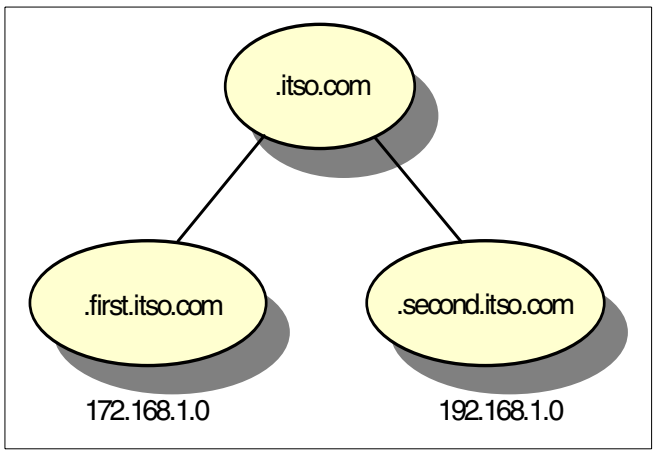

Figure 183. Domain setup

9.5.1 Setting up the master DNS

In this section we explain how we set up our master DNS server for the lab network setup. Before you start setting up the DNS server you need to check if the BIND package is installed. You can follow the instructions in 5.2,

"Setting up the Samba server" on page 155, but instead of looking for the SAMBA package, you need to search for the BIND package. After you have verified that the BIND package is installed, you need to set up your system so DNS will start automatically when the Linux server is started. To do this follow the instructions in 5.2.3, "Starting Samba as startup service" on page 166, but instead of selecting **SMB server process (samba)**, select **Internet domain name server (named)**. After you have enabled DNS to run as a service it will start automatically when the Linux server boots up.

To set up the master DNS for ".itso.com" domain follow these steps:

1. Create the /etc/named.conf file with the following entries:

```
options {
directory "/var/named";
};
zone "." {
type hint;
file "root.hint";
};
zone "localhost" {
type master;
file "localhost.zone";
};
zone "0.0.127.in-addr.arpa" {
type master;
file "127.0.0";
};
zone "itso.com" {
notify no;
type master;
file "itso.com";
};
zone "1.168.172.in-addr.arpa" {
notify no;
type master;
file "172.168.1";
};
zone "1.168.192.in-addr.arpa" {
notify no;
type master;
```

   As you can see we defined the zone file for the "itso.com" domain and the zone files for reverse address resolution for local host, network 172.168.1.0 and network 192.168.1.0.

2. Create the directory /var/named.

3. Create the zone file /var/named/itso.com with the following entries:

```
@       IN      SOA     x230.itso.com. root.itso.com.  (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
;
 NS x230 ; Name Server
 MX 10 mail ; Mail Server
;
x230 A 172.168.1.1
mail CNAME x230
first A 172.168.1.1
second A 192.168.1.1
second.itso.com. 86400 IN NS x340.second.itso.com.
x340.second.itso.com. 86400 IN A  192.168.1.2
first.itso.com. 86400 IN NS x220.first.itso.com.
```

We specified in this file that all requests for “first.itso.com” and “second.itso.com” go to the corresponding DNS servers in these domains.

4. Create the zone file /var/named/172.168.1 with the following entries:

```
@       IN      SOA     x230.itso.com. root.itso.com.  (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
 NS x230.itso.com. ; Name Server
1 PTR x230.itso.com.
2.1.168.172.in-addr.arpa. IN CNAME 2.0-255.0.168.172.in-addr.arpa.
3.1.168.172.in-addr.arpa. IN CNAME 3.0-255.0.168.172.in-addr.arpa.
;
0-255.0.168.172.in-addr.arpa.  86400 IN NS  x220.first.itso.com.
```

As you can see all requests for the reverse address resolution of the 172.168.1.0 network are passed on to the DNS in the “first.itso.com” domain. So when the DNS server gets a request for IP address in the network 172.168.1.0-255 it will pass this request to the DNS server that is serving this network. In our example, this is the DNS server for “first.itso.com” domain.

5. Create the zone file /var/named/192.168.1 with the following entries:

```
@       IN      SOA     x230.itso.com. root.itso.com.  (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
 NS x230.itso.com. ; Name Server
1 PTR x230.itso.com.
2.1.168.192.in-addr.arpa. IN CNAME 2.0-255.0.168.192.in-addr.arpa.
3.1.168.192.in-addr.arpa. IN CNAME 3.0-255.0.168.192.in-addr.arpa.
;
0-255.0.168.192.in-addr.arpa.  86400 IN NS x340.second.itso.com.
```

As you can see, all requests for the reverse address resolution of the 192.168.1.0 network are passed on to the DNS in the “second.itso.com” domain. So when the DNS server gets a request for an IP address in the network 192.168.1.0-255, it will pass this request to the DNS server that is serving this network. In our example this is the DNS server for the “second.itso.com” domain.

6. Create the zone file /var/named/127.0.0 with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
 NS localhost.
```

7. Create the zone file /var/named/localhost.zone with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
 NS localhost.
```

8. You need to set up the DNS client so it will point to the DNS server running on the server. You need to specify the address of the DNS server to be 127.0.0.1, following the procedure described in 3.19, “Name resolution settings” on page 124.

9. Using the Webmin configuration tool you need to set up the domain name to “itso.com”.

10. Your server is ready to be powered on. To start the server without restarting the operating system, which is the case in another very popular operating system, execute the command:

```
/etc/rc.d/init/named start
```

Congratulations! You have just set up a fully functional DNS server. Get ready for more excitement when you will set up the DNS servers for the subdomains.

9.5.2 Setting up the DNS for the first subdomain

Before you start configuring DNS, you need to check if the DNS server is installed. You can do this following the instructions in 9.5.1, “Setting up the master DNS” on page 216.

After you checked all prerequisites, follow these steps to set up DNS for the “first.itso.com” domain:

1. Create the /etc/named.conf file with the following entries:

```
options {
directory "/var/named";
forward only;
forwarders {172.168.1.1;};
};
zone "." {
type hint;
file "root.hint";
};
};
zone "localhost" {
type master;
file "localhost.zone";
};
zone "0.0.127.in-addr.arpa" {
type master;
file "127.0.0";
};
};
zone "first.itso.com" {
notify no;
type master;
file "first.itso.com";
};
zone "1.168.172.in-addr.arpa" {
notify no;
type master;
file "172.168.1";
};
zone "0-255.1.168.172.in-addr.arpa" {
type master;
```

As you can see we defined the zone for the root domain "." and the zone for "first.itso.com" domain. We also defined zones for reverse address resolution for local network 172.168.1.0 and for network 172.168.1.0-255, which serve the requests from the root server. You can see the zone files for network 172.168.1.0 and network 172.168.1.0-255 are the same. We need two definitions because the first one (1.168.172) is for local requests and the second (0-255.1.168.172) is for resolving requests from the master server in case someone else requests reverse resolution in this network (172.168.1.0) from the root server. The master server will ask the server serving this network (172.168.1.0) for the information (in our example, the server for the "first.itso.com" domain). Refer to 9.5.1, "Setting up the master DNS" on page 216 to see how the root server setup is done to pass requests to servers in subdomains. As you can see in the options section, we set up forwarding, because the server is on a private network and it cannot reach the root servers. Therefore, any requests that are not

for the “first.itso.com” domain will be passed to the master server of the “itso.com” domain.

2. Create the directory /var/named.
3. Create the zone file /var/named/first.itso.com with the following entries:

```
@       IN      SOA x220.first.itso.com. root.first.itso.com. (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
;
 NS x220 ; Name Server
 MX 10 mail ; Mail Server
;
x220 A 172.168.1.2
client A 172.168.1.3
mail CNAME x220
```

In this file we create definitions for all the computers in network 172.168.1.0.

4. Create the zone file /var/named/172.168.1 with the following entries:

```
@       IN      SOA x220.first.itso.com. root.first.itso.com. (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
 NS x220.first.itso.com. ; Name Server
1 PTR x230.itso.com.
2 PTR x220.first.itso.com.
3 PTR client.first.itso.com.
```

In this file we define reverse address resolution for network 172.168.1.0.

5. Create the zone file /var/named/127.0.0 with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                        1997022700 ; Serial
                                        28800      ; Refresh
                                        14400      ; Retry
                                        3600000    ; Expire
                                        86400 )    ; Minimum
 NS 127.0.0.1
```

6. Create the zone file /var/named/localhost.zone with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
 NS localhost.
```

7. You need to set up the DNS client so it will point to the DNS server running on the server. You need to specify the address of the DNS server to be 127.0.0.1, following the procedure described in 3.19, “Name resolution settings” on page 124.
8. Using the Webmin configuration tool you need to set up the domain name to “first.itso.com”.
9. Start the server with the command:

   ```
   /etc/rc.d/init/named start
   ```

9.5.3 Setting up the DNS for the second subdomain

Before you start configuring DNS, you need to check if the DNS server is installed. You can do this by following the instructions in 9.5.1, “Setting up the master DNS” on page 216.

After you have checked all prerequisites, follow these steps to set up the DNS for the “second.itso.com” domain:

1. Create the /etc/named.conf file with the following entries:

```
options {
directory "/var/named";
forward only;
forwarders {192.168.1.1;};
};
zone "." {
type hint;
file "root.hint";
};
};
zone "localhost" {
type master;
file "localhost.zone";
};
zone "0.0.127.in-addr.arpa" {
type master;
file "127.0.0";
};
};
zone "second.itso.com" {
notify no;
type master;
file "second.itso.com";
};
zone "1.168.192.in-addr.arpa" {
notify no;
type master;
file "192.168.1";
};
zone "0-255.1.168.192.in-addr.arpa" {
type master;
file "192.168.1";
```

As you can see we defined the zone for the root domain "." and the zone for "second.itso.com" domain. We also defined zones for reverse address resolution for local, network 192.168.1.0 and for network 192.168.1.0-255, which serves the requests from the root server. You can see the zone files for network 192.168.1.0 and network 192.168.1.0-255 are the same. We need two definitions because the first one (1.168.192) is for local requests and the second (0-255.1.192.168) is for resolving requests from the root server in case someone else requests reverse resolution in this network (192.168.1.0) from the root server. The root server will ask the server serving this network (192.168.1.0) for the information (in our example, the server for the "first.itso" domain). Refer to 9.5.1, "Setting up the master DNS" on page 216 to see how root server setup is done to pass requests to servers in subdomains. As you can see in the options part we se tup forwarding,because the server is on a private network and it cannot reach

the root servers. Therefore, any requests that are not for the “second.itso.com” domain will be passed to the master server of the “itso.com” domain.

2. Create the directory /var/named.
3. Create the zone file /var/named/second.itso.com with the following entries:

```
@       IN SOA x340.second.itso.com root.second.itso.com. (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
;
 NS x340 ; Name Server
 MX 10 mail ; Mail Server
;
x340 A 192.168.1.2
mail CNAME x340
```

In this file we create definitions for all computers in network 192.168.1.0.

4. Create the zone file /var/named/192.168.1 with the following entries:

```
@       IN SOA x340.second.itso.com root.second.itso.com. (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
 NS x340.second.itso.com. ; Name Server
1 PTR x230.itso.com.
2 PTR x340.second.itso.com.
3 PTR client.second.itso.com.
```

In this file we define reverse address resolution for network 192.168.1.0.

5. Create the zone file /var/named/127.0.0 with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
 NS 127.0.0.1
```

6. Create the zone file /var/named/localhost.zone with the following entries:

```
@       IN      SOA     localhost. root.localhost.  (
                                      1997022700 ; Serial
                                      28800      ; Refresh
                                      14400      ; Retry
                                      3600000    ; Expire
                                      86400 )    ; Minimum
 NS localhost.
```

7. You need to set up the DNS client so it will point to the DNS server running on the server. You need to specify the address of the DNS server to be 127.0.0.1, following the procedure described in 3.19, “Name resolution settings” on page 124.
8. Using the Webmin configuration tool you need to set up the domain name to “second.itso.com”.
9. Start the server with the command:

   ```
   /etc/rc.d/init/named start
   ```

You now have three DNS servers running. The network is ready for the setup of the mail server.

9.5.4 Setting up sendmail

All documentation on sendmail will tell you that the sendmail configuration file /etc/sendmail.cf is a nightmare for a network administrator. This is not entirely true; when you do not need any special features offered by sendmail, the setup is fairly easy. You just need to modify the generic macro files slightly and recreate the new sendmail.cf file with the “m4” macro processor. In this section we explain how to set up sendmail for handling mail in its own domain. Caldera OpenLinux comes with a generic macro file for sendmail. The file is located in the /usr/share/sendmail/cf/cf directory and is called generic-col2.2.mc. This file just needs a little modification to become a working file on your server. Follow these steps to set up your mail servers:

1. Make a copy of Caldera OpenLinux generic macro file with the command:

   ```
   cp /usr/share/sendmail/cf/cf/generic-col2.2.mc
   /usr/share/sendmail/cf/cf/mydomain.mc
   ```

2. Add the following lines to mydomain.mc file:

   a. For the server in “first.itso.com” domain:

```
dnl ###################################################################
dnl # Definitions for sample domain
dnl # we define PSEUDONYMS, DEFAULT_HOST
define(`PSEUDONYMS', `x220.first.itso.com' `first.itso.com')
define(`DEFAULT_HOST', 'x220.first.itso')
```

b. For the server in “second.itso.com” domain:

```
dnl ###################################################################
dnl # Definitions for sample domain
dnl # we define PSEUDONYMS, DEFAULT_HOST
define(`PSEUDONYMS', `x340.second.itso.com' `second.itso.com')
define(`DEFAULT_HOST', 'x340.second.itso.com')
dnl ###################################################################
```

3. Create a new sendmail.cf file with your domain file by executing the command:

```
/usr/bin/m4 /usr/share/sendmail/cf/m4/cf.m4
/usr/share/sendmail/cf/cf/mydomain.cf
/usr/share/sendmail/cf/feature/relay_entire_domain.m4 >
/etc/sendmail.cf
```

This can be a lot easier if your current directory is /usr/share/sendmail/cf:

```
/usr/bin/m4 m4/cf.m4 cf/mydomain.cf feature/relay_entire_domain.m4 >
/etc/sendmail.cf
```

As you can see, we used two more files in order to create the configuration file:

a. The cf.m4 file must be used. Otherwise files will not be parsed correctly.

b. The relay_entire_domain.m4 file is used to enable clients that are accessing the mail server with remote clients, to send mail through this server.

4. Modify the /etc/sendmail.cw file:

a. For the server in “first.itso.com” domain add the line:

```
first.itso.com
```

a. For the server in “second.itso” domain add the line:

```
second.itso.com
```

The /etc/sendmail.cw includes all aliases for your mail server. You need to include the domain name; otherwise, mail will be undeliverable.

5. Start the server with the command:

```
/etc/rc.d/init.d/mta start
```

If you want sendmail to start automatically when the server is started follow the instructions in 5.2.3, “Starting Samba as startup service” on page 166, but instead of selecting SMB server process (samba) you select **Mail Transfer Agent**.

You need to execute all the previous instructions on the servers in both domains (first.itso and second.itso) if you want to send mail from one domain to another.

9.5.4.1 Configure sendmail for mail routing

By default sendmail can deliver mail to a defined user if it can reach the mail server for the user’s domain. So, for example, if you are on an internal network and your sendmail server does not have direct connection to the Internet you can configure sendmail to route mail through another reachable mail server that is connected to the Internet. To do this you need to enter the appropriate values into the /etc/sendmail/mailertable file. For example, if you want to route the mail for `otherdomain.com` through `reachableserver`, then your /etc/sendmail/mailertable file will look similar to this:

```
#
# /etc/mailertable
#
# This file can be used to fine-tune sendmail's email routing.
# After you change this file, you must rebuild the DB file using
# /usr/sbin/makemap hash /etc/sendmail/mailertable < /etc/sendmail/mailer
#
otherdomain. smtp:reachableserver
```

This means that all mail for `otherdomain` will be routed to the `reachableserver` mail server with the SMTP protocol. The domain name must always have a “.” dot at the end. With this setup you can route mail from server to server until it reaches its destination.

Whenever you modify the /etc/sendmail/mailertable file, you need to rebuild the /etc/sendmail/mailertable.db file, which is really the file used by sendmail to perform routing tasks. You can do this by executing the command:

```
/usr/sbin/makemap hash /etc/sendmail/mailertable <
/etc/sendmail/mailertable
```

If you want all your mail to be routed to another server, you may do this by appending the following line to the /etc/sendmail/mailertable file:

```
smtp:smartserver
```

Where `smartserver` is the mail server that will handle all mail from your sendmail server.

9.5.5 Configuring sendmail using Webmin

Caldera OpenLinux eServer Version 2.3 includes a Web-enabled tool for system administration called Webmin. In this section we will show you how to set up sendmail using this tool. You can see how to start the Webmin tool in 3.22, “System administration using Webmin” on page 129.

Before you start configuring the sendmail with Webmin you need to create the original sendmail.cf file by executing the command:

```
/usr/bin/m4 /usr/share/sendmail/cf/m4/cf.m4
/usr/share/sendmail/cf/cf/generic-col2.2.mc
/usr/share/sendmail/cf/feature/relay_entire_domain.m4 >
/etc/sendmail.cf
```

As you can see we used default settings only, because we will customize all the options later. In our example we will show you the customization of the sendmail server in the “second.itso.com” domain. To customize the sendmail server follow these steps:

1. In the main Webmin window select the **Servers** tab. You will see a window similar to Figure 184.

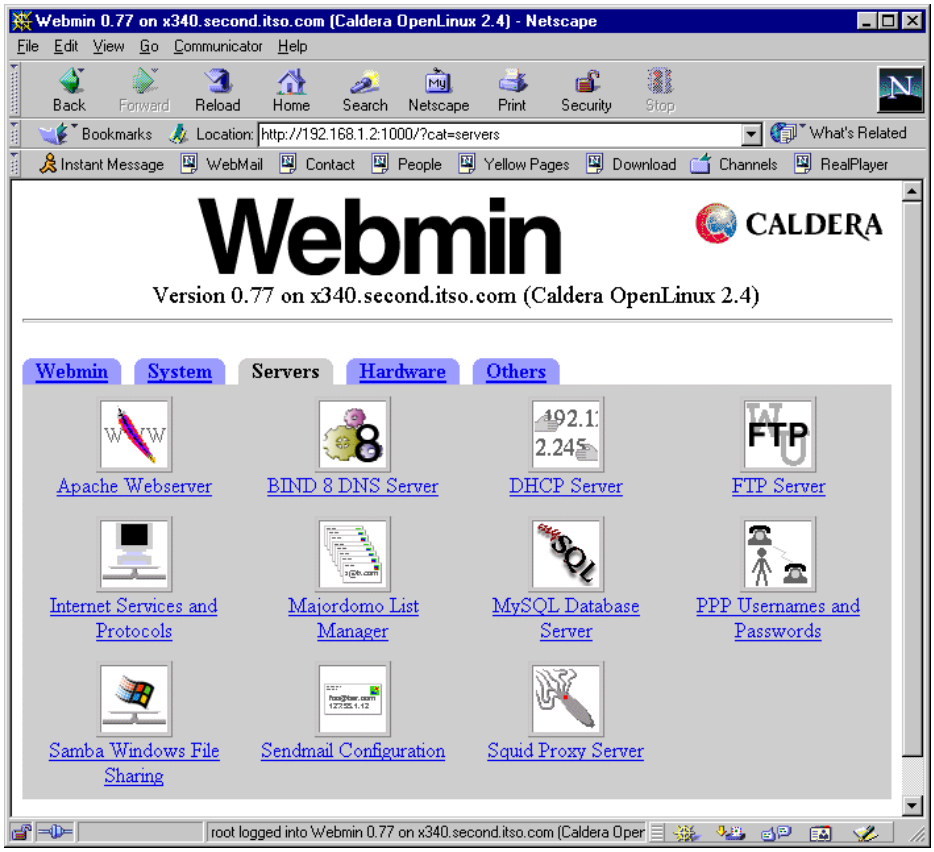

Figure 184. Webmin servers

2. Select the **Sendmail Configuration** icon. You will see a window similar to Figure 185.

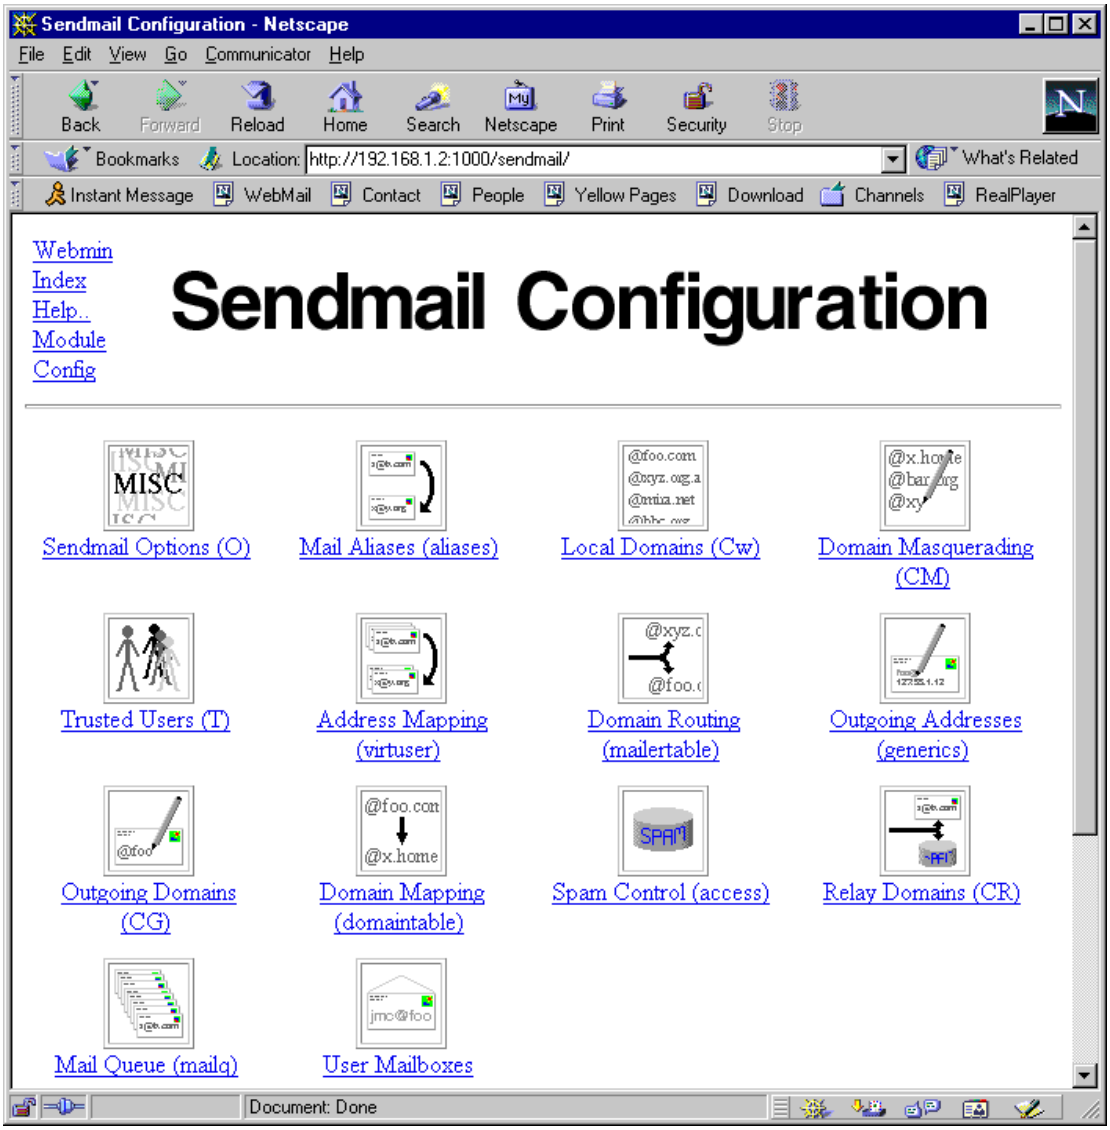

Figure 185. sendmail configuration

3. First we will set up Domain Masquerading. This is useful if you do not want mail to be represented by the full hostname, such as “col@client.second.itso.com”, but just by the domain name “col@second.itso.com”. Select **Domain Masquerading** and you will see a window similar to Figure 186.

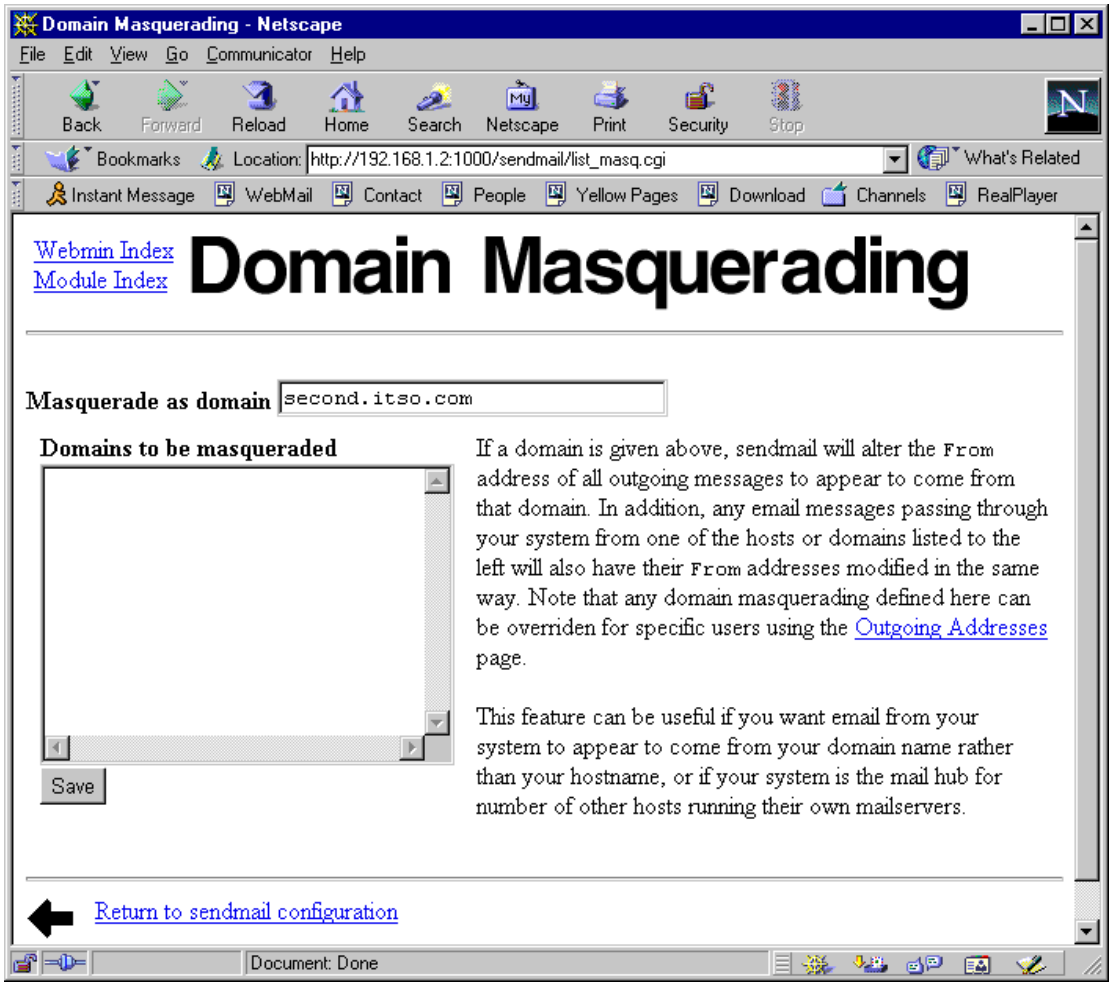

Figure 186. Domain Masquerading

In the **Masquerade as domain** field, type the domain name you want to use for masquerading. In our example, this is “second.itso.com”. Click **Save** to save your changes and then return back to the previous window.

4. In this step we will specify the domain names from which our mail server will accept the mail. Select the **Local Domains** and you will see a window similar to Figure 187.

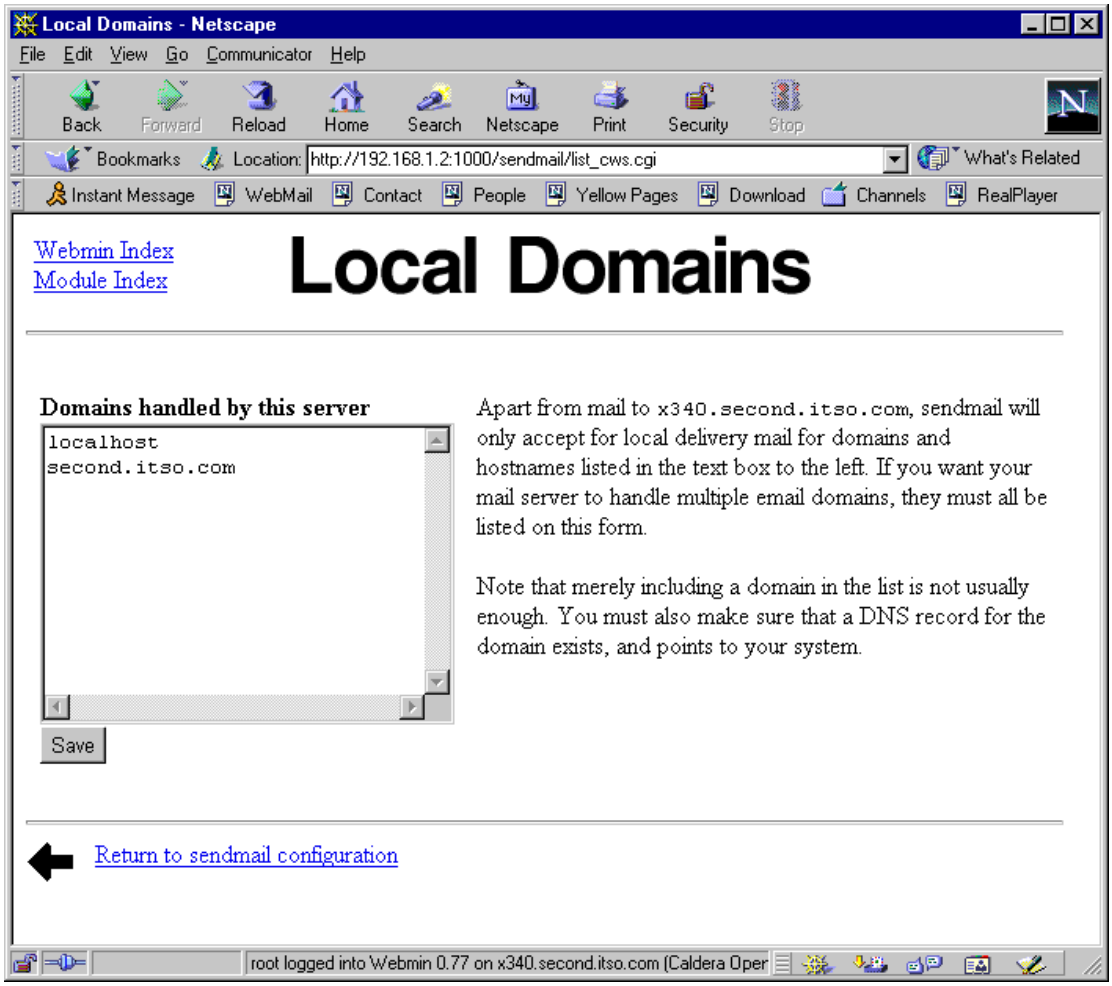

Figure 187. Local Domains

Type in the domain names from which your mail server will receive mail. In our example this is “second.itso.com”.

After this you need to restart the mail server with the command:

```
/etc/rc.d/init.d/mta reload
```

You need to repeat this also for the mail server in the “first.itso.com” domain.

9.5.6 Setting up the mail client

Each mail user needs a user name/password on the server that is running mail server (in our case sendmail). After the user has a user name/password, he can reach his mailbox remotely. In this section we show you how to set up the Netscape mail setting for sending and receiving mail.

Note

The POP3 server must be installed and running before you can set up the client.

Follow these steps to set up the Netscape mail properties:

1. Start Netscape.
2. Select **Edit > Preferences > Mail & Newsgroups > Identity** and you will see a window similar to Figure 188.

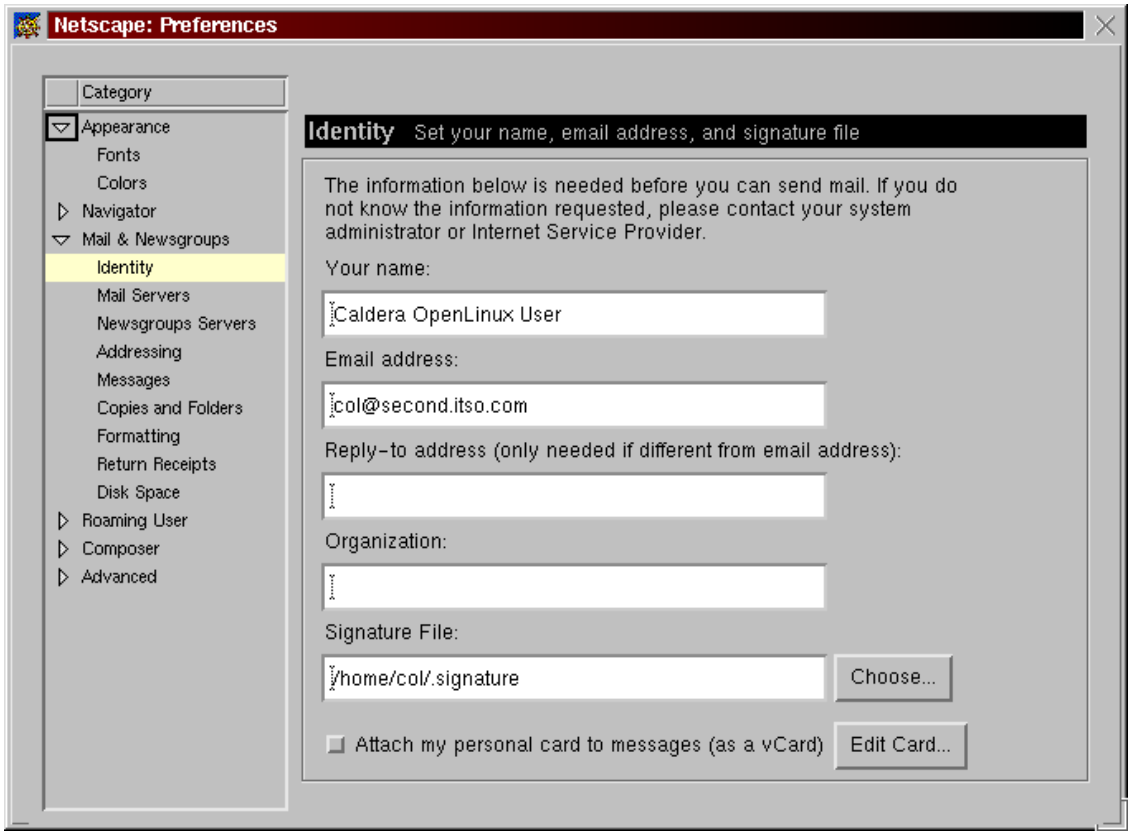

Figure 188. Setting the identity settings

3. Type in the required values and select **Mail Servers.** You will see a window similar to Figure 189.

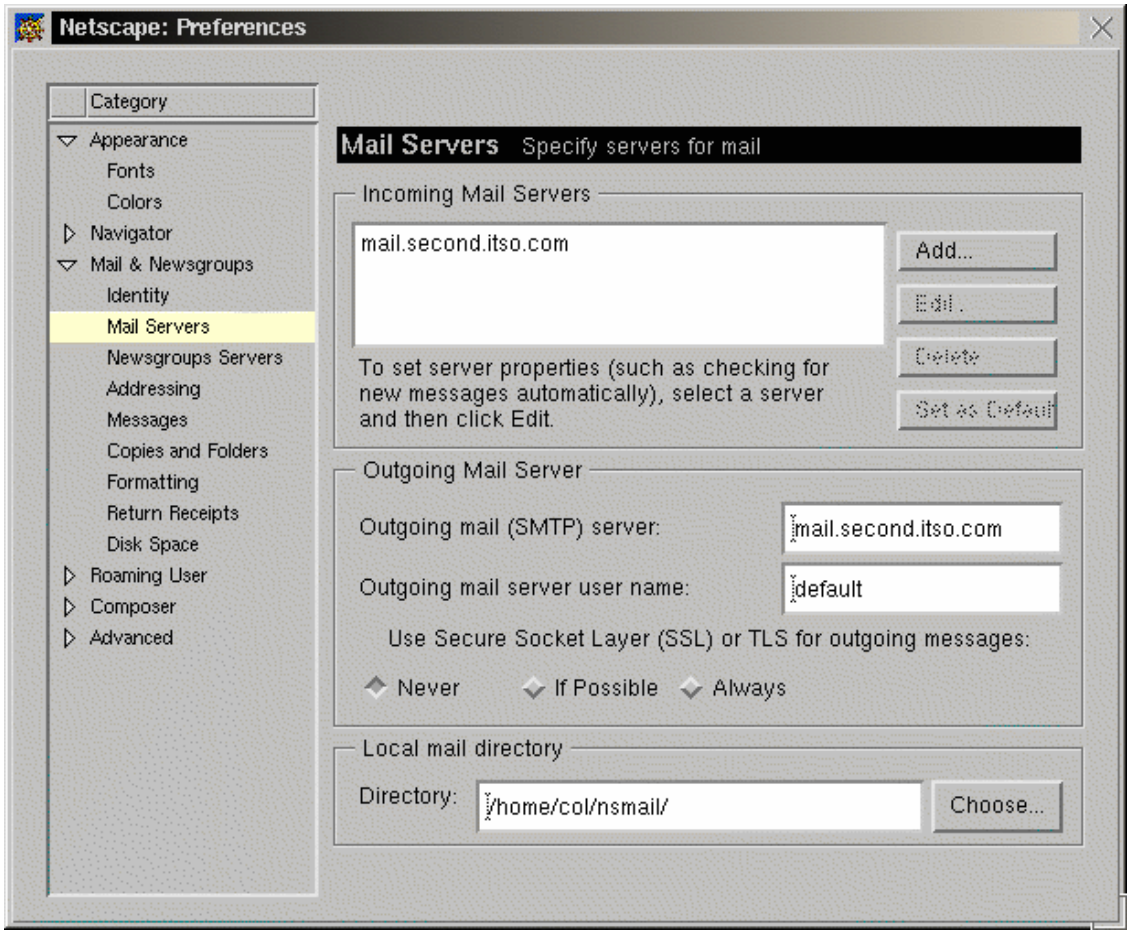

Figure 189. Setting the mail servers

4. In the Outgoing mail (SMTP) server field, type your mail server address (in our example, “mail.second.itso.com”).
5. In the Incoming Mail Server section, select the current server and click **Edit**, and you will see a window similar to Figure 190.

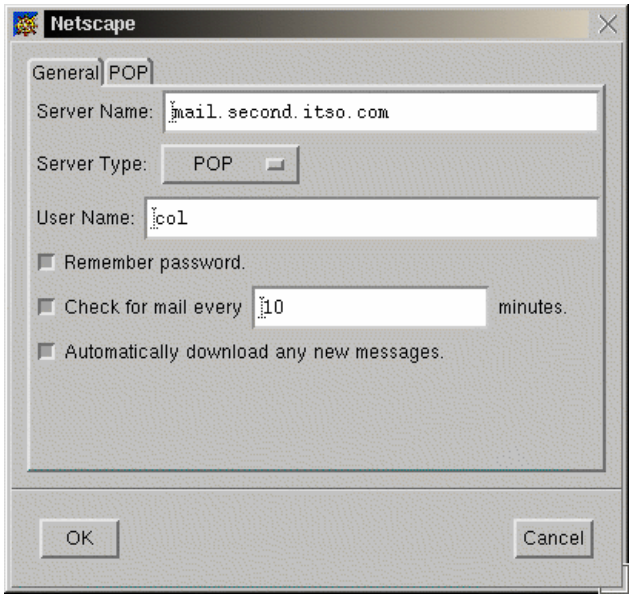

Figure 190. Setting POP3 server

6. In the Server Name field, type in the address of your mail server.
7. In the User Name field, type in your user name on the mail server. You can also configure some other options that will affect your mail reading. Click **OK** to continue.
8. When you are back in the Preferences window, click **OK** to store your new mail settings.

Now you are ready to send mail to all the users in the “first.itso.com” and “second.itso.com” domain.

9.6 Sources of additional information

You can find more information on the official Web site of the sendmail project:

```
http://www.sendmail.org
```

And there are always good how-to documents at the Linux documentation project Web site:

```
http://www.linuxdoc.org/
```

Chapter 10. 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.

10.1 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.

10.1.1 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.

10.1.2 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.

10.2 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

10.3 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=ldap1, ou=people, dc=ibm, dc=com
uid: ldap1
cn: ldap1
objectclass: account
objectclass: posixaccount
objectclass: top
objectclass: shadowaccount
userpassword:test
shadowLastchange: 11263
shadowMax: 99999
shadowWarning: 7
loginShell: /bin/bash
uidNumber: 601
gidNumber: 601
homeDirectory: /home/ldap1
```

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(domain name or dn), ibm(dn), and people(organizational unit).

Within the organizational unit, we will store the users and their corresponding authentication information.The second group of entries is to add the user ldap1 and the attributes that are necessary for authentication.

10.4 Implementation on Linux

You can download the latest stable version of OpenLDAP from:

```
ftp.openldap.org/pub/openldap
```

However, we will be using the RPMs that come with the Caldera OpenLinux eServer 2.3 distribution. Install or verify installation of the following RPMs:

openldap-1.2.7-2

nss_ldap-85-1

10.4.1 Slapd.conf

Now edit the /etc/ldap/slapd.conf file. Replace ibm.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/ldap/slapd.at.conf
include         /etc/ldap/slapd.oc.conf
schemacheck     off
#referral       ldap://ldap.itd.umich.edu

pidfile         /var/run/slapd.pid
argsfile        /var/run/slapd.args

########################################################################
# ldbm database definitions
########################################################################

database        ldbm
suffix          "ou=people, dc=ibm, dc=com"
#suffix         "o=Your Organization Name, c=US"
directory       /var/ldap
rootdn          "cn=admin, ou=people, dc=ibm, dc=com"
#rootdn         "cn=root, o=Your Organization Name, c=US"
rootpw          secret
# cleartext passwords, especially for the rootdn, should
# be avoid.  See slapd.conf(5) for details.
```

10.4.2 ldap.conf

When the nss_ldap rpm was installed it created the file /etc/ldap.conf. Edit the /etc/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.

```
# This is the configuration file for the LDAP nameservice
# switch library and the LDAP PAM module.
# To contact the developers, mail support@padl.com.
# If the host and base aren't here, then the DNS RR
# _ldap._tcp.<defaultdomain>. will be resolved. <defaultdomain>
# will be mapped to a distinguished name and the target host
```

```
# will be used as the server.

# Your LDAP server. Must be resolvable without using LDAP.
host 127.0.0.1

# The distinguished name of the search base.
base ou=people,dc=ibm,dc=com

# The LDAP version to use (defaults to 2)
#ldap_version 3
# The distinguished name to bind to the server with.
# Optional: default is to bind anonymously.
#binddn cn=manager,dc=padl,dc=com
# The credentials to bind with.
# Optional: default is no credential.
#bindpw secret
# The port.
# Optional: default is 389.
#port 389
# The search scope.
#scope sub
#scope one
#scope base
# The following options are specific to nss_ldap.
# The hashing algorith your libc uses.
# Optional: default is des
#crypt md5
#crypt sha
#crypt des
# The following options are specific to pam_ldap.
# Filter to AND with uid=%s
#pam_filter objectclass=account
# The user ID attribute (defaults to uid)
#pam_login_attribute uid
# Search the root DSE for the password policy (works
# with Netscape Directory Server)
#pam_lookup_policy yes
# Group to enforce membership of
#pam_groupdn cn=PAM,ou=Groups,dc=padl,dc=com
# Group member attribute
#pam_member_attribute uniquemember
# Hash password locally; required for University of
# Michigan LDAP server, and works with Netscape
# Directory Server if you're using the UNIX-Crypt
# hash mechanism and not using the NT Synchronization
# service.
#pam_crypt local
```

10.4.3 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 Service 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:

```
passwd: compat ldap
```

10.4.4 /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 the user's password.
- password: authentication update mechanism. Example: Prompting the user to enter a new password when their 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:

```
auth       required     /lib/security/pam_securetty.so
auth       required     /lib/security/pam_nologin.so
auth       sufficient   /lib/security/pam_ldap.so
auth       required     /lib/security/pam_unix_auth.so try_first_pass
account    sufficient   /lib/security/pam_ldap.so
account    required     /lib/security/pam_unix_acct.so
password   required     /lib/security/pam_cracklib.so
password   required     /lib/security/pam_ldap.so
password   required     /lib/security/pam_pwdb.so use_first_pass
session    required     /lib/security/pam_unix_session.so
session    optional     /lib/security/pam_console.so
```

For more information on PAM, see the PAM Administrators guide at:
`http://www.kernel.org/pub/linux/libs/pam/Linux-PAM-html/pam.html`

10.4.5 Starting OpenLDAP

To start slapd, type `/usr/sbin/slapd` (-d 255 for debugging information).

With slapd successfully running, we now need to load the initial database. Create an LDIF file like the one on page 241. Replace `ldap1` 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 secret
```

10.4.6 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/ldap1
chown 500:500 /home/ldap1
```

Now try to log in as ldap1 with the password test.

If the authentication failed, check the configuration files for typos. It can also be helpful to run slapd in debug mode to watch the output of the failed authentication.

Once you have logged in successfully, either remove the test user or at the very least create a properly encrypted password.

Chapter 11. 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.

11.1 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 an NIS-based environment. NISGINA provides an 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/
```

11.2 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

11.3 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 box. 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.

Packages that need to be installed for an NIS client and the NIS server:

- nis-client-2.0-8
- nis-server-2.0-8

The /etc/nsswitch.conf file determines the order of lookups performed.

Sample /etc/nsswitch.conf file:

```
# /etc/nsswitch.conf
#
# Name Service Switch configuration file.
#
passwd : compat
shadow : compat
group : compat
#
hosts : files dns nis
networks : nis files dns
#
ethers : nis files
protocols : nis files
rpc : nis files
services : nis files
```

Note

You should remove NIS references from all the options where you do not use the NIS server for coordination.

11.3.1 NIS Server

The configuration files for the NIS server in Caldera OpenLinux are located in the directory /etc/nis. To set up an NIS server on your Caldera OpenLinux server, follow these steps:

1. Decide on an NIS domain name. This name does not need to be equal to the domain name of the server because it only has to be unique among adjacent domains, not worldwide.
2. Create a directory /etc/nis/nisdomainname. In our example we used domain name “nis.com”. We created a directory by executing this command:

   ```
   mkdir /etc/nis/nis.com
   ```

3. You can define the NIS domain using COAS. Click the **K** sign on the panel, select **COAS -> Network -> TCP/IP -> NIS**. See Figure 191.

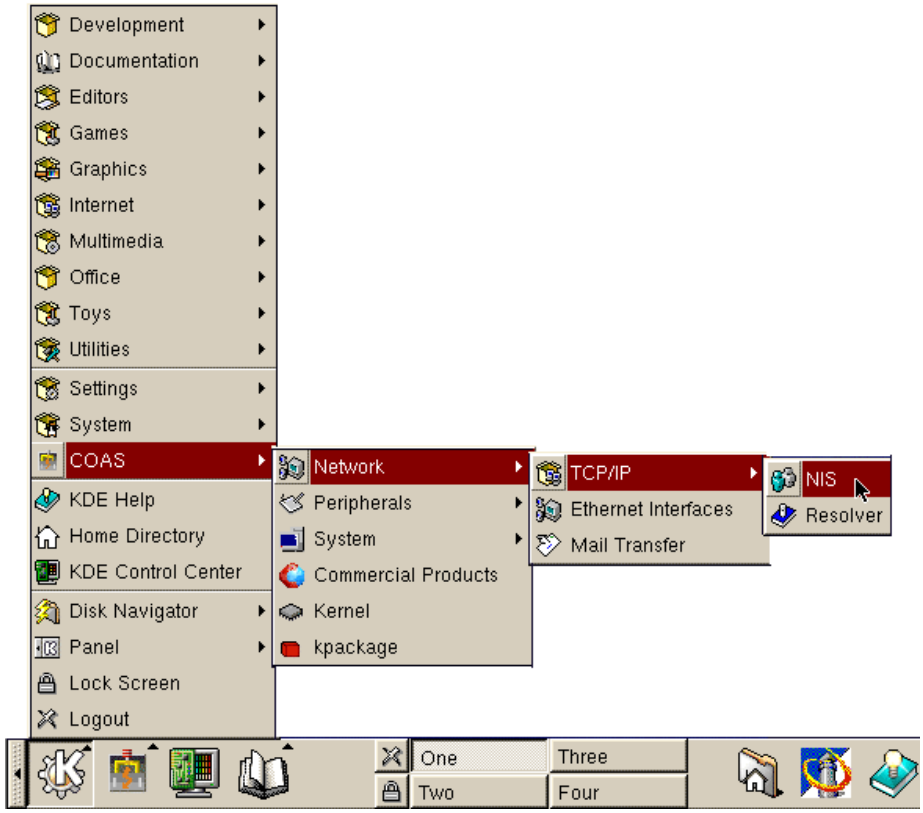

Figure 191. Starting NIS settings

After the NIS configuration is started you will see a window similar to Figure 192.

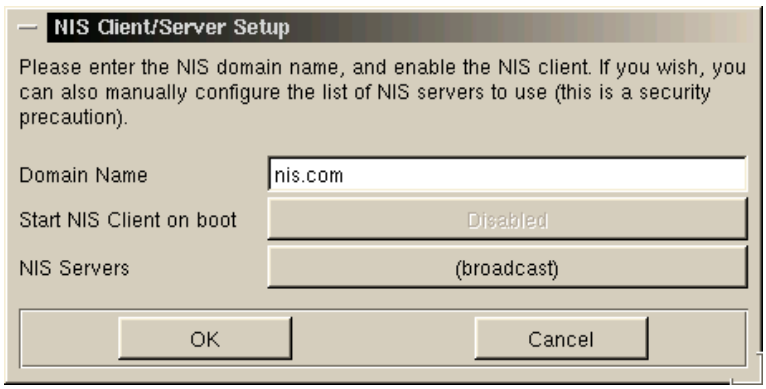

Figure 192. NIS properties

4. In the Domain Name field, type in your domain name. In our example type in `nis.com`.
5. Click **OK** to continue, **Save** on the next window and to finish select **Done** on the last window or wait until the window closes automatically.
6. Create symbolic links of all files you want to have shared in the local network into /etc/nis/nisdomainname. In our example we want to have password and group files distributed, so we created symbolic links for these files:

   ```
   ln -s /etc/passwd /etc/nis/nis.com/passwd
   ln -s /etc/shadow /etc/nis/nis.com/shadow
   ln -s /etc/group /etc/nis/nis.com/group
   ```

7. Copy /etc/nis/.nisupdate.conf.sample to the directory for your domain, in our example /etc/nis/nis.com, with the filename changed to .nisupdate.conf, by executing the command:

   ```
   cp /etc/nis/.nisupdate.conf.sample /etc/nis/nis.com/.nisupdate.conf
   ```

   This file is used to update or create maps used by the NIS server. Edit this file for the maps you want.
8. Run `/etc/nis/nis_update`. You will see output similar to Figure 193. By executing this command all NIS domain directories are scanned and the files for the NIS server are created in the directory /var/nis.

```
[root@x230 nis]# /etc/nis/nis_update
Processing domain nis.com
Updating nis.com/passwd
Updating nis.com/group
```

Figure 193. After executing nis_update

Do this every time one of the files used by NIS server changes. If you examine this script you can find examples of which files to put into the NIS environment as well.

9. Start the NIS server by executing the command:

   ```
   /etc/rc.d/init/nis-server start
   ```

If you want the NIS server to start automatically when the server is started, follow the instructions in 5.2.3, “Starting Samba as startup service” on page 166, but instead of selecting SMB server process (samba), select **NIS--Network Information Service (server part)**.

Take a look at /var/yp/securenets. This file defines the access rights to your NIS server. By default it is set to give access to everyone. Change it accordingly (see man securenets).

To test our NIS server 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
```

To test our NIS master server, we need to set up a client to run ypbind, which is a client for NIS servers.

11.3.2 NIS Client

The client can be set up with COAS. Follow these steps to set up the NIS client:

1. Start the NIS configuration as we described in 11.3.1, “NIS Server” on page 249. When you start the NIS configuration you will see a window similar to Figure 194.

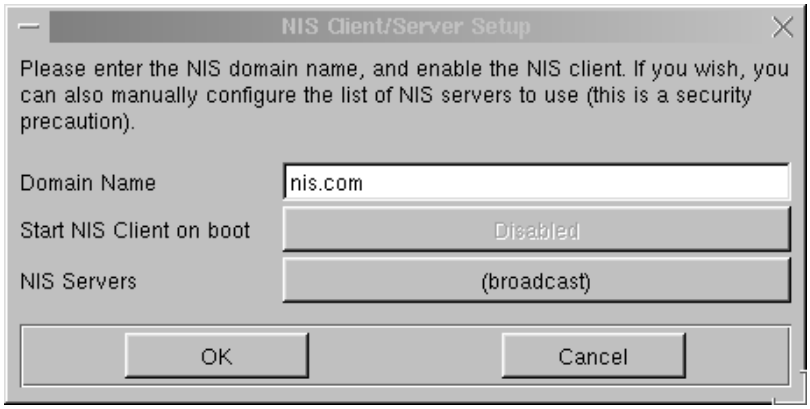

Figure 194. Configuring the NIS client

2. In the Domain Name field, type in the NIS domain name.
3. Click the **NIS Servers** button. You will see a window similar to Figure 195.

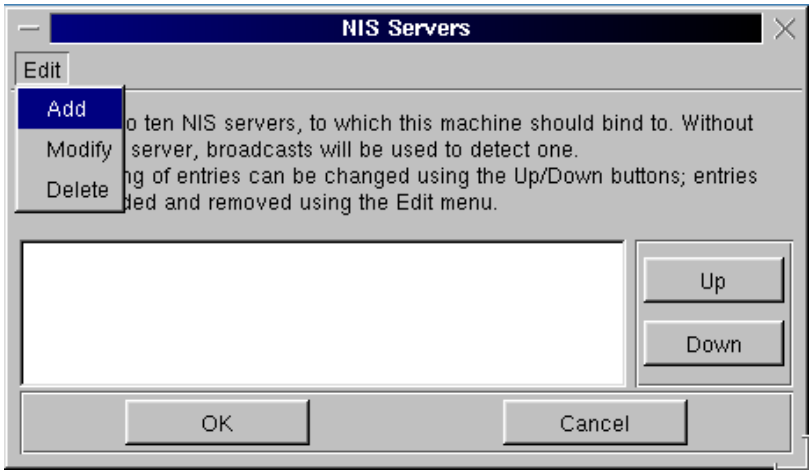

Figure 195. Adding NIS server

4. In Edit menu select **Add.** You will see a window similar to Figure 196.

Figure 196. Specifying the address of NIS server

5. Type in the hostname or the IP address of your NIS server. Click **OK** to continue. You will see a window similar to Figure 197.

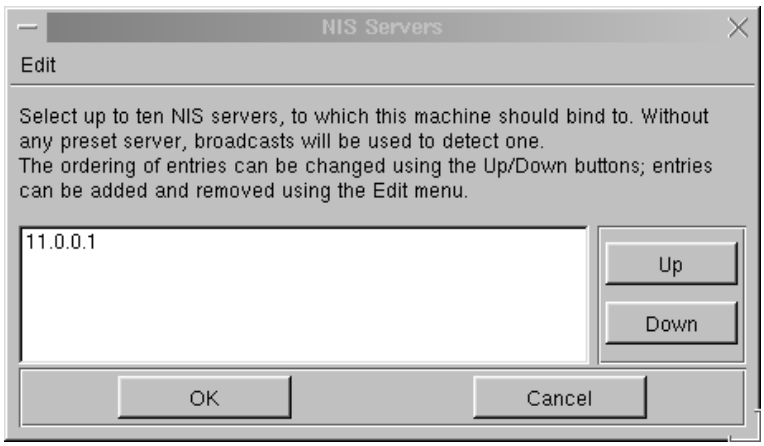

Figure 197. After defining the NIS server

6. Click **OK** to continue. Select **Save** at the next window, and you will see that your NIS client will be restarted. Click **Done** to finish the setup or wait until the window closes automatically.
7. Start the NIS client by executing the command:

   ```
   /etc/rc.d/init/nis-client start
   ```

If you want the NIS client to start automatically when the server is started follow the instructions in 5.2.3, “Starting Samba as startup service” on

page 166, but instead of selecting SMB server process (samba), select **NIS--Network Information Service (client part)**.

To test our NIS configuration we will use the `ypcat` command:

```
ypcat passwd
```

You should see output similar to Figure 198.

```
[root@x220 /root]# ypcat passwd
gopher:*:13:30:gopher:/usr/lib/gopher-data
bin:*:1:1:bin:/bin
col:awfZA7yDE5TH6:500:100:Caldera Systems OpenLinux User:/home/col:/bin/bash
ftp:*:14:50:FTP User:/home/ftp
shutdown:*:6:11:shutdown:/sbin:/sbin/shutdown
nobody:*:65534:65534:Nobody:/:/bin/false
lp:*:4:7:lp:/var/spool/lpd
halt:*:7:0:halt:/sbin:/sbin/halt
daemon:*:2:2:daemon:/sbin
postgres:*:17:17:Postgres User:/home/postgres:/bin/bash
root:awfZA7yDE5TH6:0:0:root:/root:/bin/bash
mysql:*:18:18:MySQL User:/var/lib/mysql:/bin/false
games:*:12:100:games:/usr/games
man:*:15:15:Manuals Owner:/
mail:*:8:12:mail:/var/spool/mail
majordom:*:16:16:Majordomo:/:/bin/false
sync:*:5:0:sync:/sbin:/bin/sync
news:*:9:13:news:/var/spool/news
ivo:wUQ0uaL7zkTf.:501:501:Ivo:/home/ivo:/bin/bash
operator:*:11:0:operator:/root
adm:*:3:4:adm:/var/adm
uucp:*:10:14:uucp:/var/spool/uucp
[root@x220 /root]#
```

Figure 198. Executing ypcat passwd

Before you can log on to the local machine with NIS, using the account defined on the NIS server you need to modify /etc/nsswitch.conf file to define that passwd will be searched in the NIS database. The modified file should look similar to the following:

```
# /etc/nsswitch.conf
#
# Name Service Switch configuration file.
#
passwd:         nis
shadow:         nis
group:          nis
#
hosts:          files nis dns
networks:       nis files dns
```

```
#
ethers:         nis files
protocols:      nis files
rpc:            nis files
services:       nis files
```

Now to really test the machine, log in to an NIS client using an account that is on the NIS master server. When you log in, you should see the following:

```
Caldera OpenLinux eServer
Version 2.3
Copyright 1996-1999 Caldera Systems, Inc.



x220.first.itso.com login: ivo
Password:
No directory /home/ivo!
Logging in with home = "/".
bash$
```

Figure 199. No home directory

Since Ivo's home directory is not defined on the client box, but only on the NIS master server, we get an error when logging in. This can be fixed by creating a home directory for Ivo on the client box if necessary. Another option would be to use NFS in conjunction with NIS to automatically mount Ivo's home directory.

11.4 Sources of additional information

For further information or troubleshooting, *The Linux NIS (YP)/NYS/NIS+ HOWTO* by Thorsten Kukuk is a good place to start. Find it at
`http://www.metalab.unc.edu/pub/Linux/docs/HOWTO/NIS-HOWTO`

Managing NFS and NIS by Hal Stern is also a good resource.

Chapter 12. 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 filesystem on your computer. The mounting process is the same as for any filesystem 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.

12.1 The NFS process

First you need to verify that the NFS utilities RPM package has been loaded. RPM 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 -ql nfs-server
```

You will see the results as in Figure 200. This gives a listing of the contents of the RPM that are installed when the NFS utilities are installed. The NFS utilities may have been installed by default when you first set up the system by choosing a package group that included it. You can also verify if the package is installed with the kpackage utility from the COAS tools as we described in 3.5, “Adding and removing software packages using kpackage” on page 93. The nfs-server package is found by clicking **Server -> Network**.

You can also install the NFS utilities manually by using the kpackage utility or by using the following commands and installing the RPM:

```
mount /mnt/cdrom
rpm -ivh /mnt/cdrom/col/install/RPMS/nfs-server-2.2beta47-2.i386.rpm
```

The NFS utilities package should then be installed. Now you should repeat the earlier `rpm` command to verify that the package is installed.

```
[root@x230 /]# rpm -ql nfs-server
/etc/rc.d/init.d/nfs
/etc/sysconfig/daemons/nfs
/usr/doc/nfs-server-2.2beta47
/usr/doc/nfs-server-2.2beta47/NEWS
/usr/doc/nfs-server-2.2beta47/README
/usr/man/man5/exports.5.gz
/usr/man/man8/mountd.8.gz
/usr/man/man8/nfsd.8.gz
/usr/man/man8/rpc.mountd.8.gz
/usr/man/man8/rpc.nfsd.8.gz
/usr/sbin/rpc.mountd
/usr/sbin/rpc.nfsd
[root@x230 /]#
```

Figure 200. Checking NFS RPMs

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.mountd — This handles the exporting of NFS filesystems. It looks in the /etc/exports file to figure out what to do with mount requests from various hosts.

rpc.nfsd — This provides the user level part of the NFS process.

rpc.rquotad — This handles quotas for access to filesystems. The quotas are based on disk usage and can be hard or soft limits.

You can verify that the rpc.nfsd, rpc.mountd, and portmap daemons are running as shown in Figure 201.

```
Starting NFS services:  mountd nfsd.
[root@x230 col]# ps ax | grep nfs
  901 ?        S      0:00 nfsd 4
  903 ?        S      0:00 nfsd 4
  904 ?        S      0:00 nfsd 4
  905 ?        S      0:00 nfsd 4
[root@x230 col]# ps ax | grep mount
  899 ?        S      0:00 mountd
  909 pts/1    S      0:00 grep mount
[root@x230 col]# ps ax | grep portmap
  512 ?        S      0:00 rpc.portmap
[root@x230 col]#
```

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:

```
/etc/rc.d/init.d/portmap start
```

Once the portmap daemon is running, you can start up the NFS daemons with the command:

```
/etc/rc.d/init.d/nfs start
```

```
[root@x230 col]# /etc/rc.d/initortnfs start
Starting NFS services:  mountd nfsd.
[root@x230 col]#
```

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 12.2, “Allowing NFS access to data” on page 260.

To stop the NFS server you use the command:

```
/etc/rc.d/init.d/nfs stop
```

The results are shown in Figure 203. You will notice that some processes are shut down, but not necessarily in the same order. The quota process is started up first because the quotas need to be established before the mounts take place. It is shut down last, so that nothing can slip past the quota process.

```
[root@x230 col]# /etc/rc.d/init.d/nfs stop
Stopping NFS services:  nfsd mountd.
[root@x230 col]#
```

Figure 203. Stopping the NFS server

You can restart the NFS process with the command:

```
/etc/rc.d/init.d/nfs restart
```

This can also be used to restart the NFS process if you have made changes to the configuration files.

12.2 Allowing NFS access to data

You can give NFS access to a filesystem 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 filesystem 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 17.

You need to be sure that the exporting server can recognize the server name.

The various options are explained in the table below.

Table 17. Access options

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 an 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 below in Figure 204.

```
 # 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 204. 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 accessible read-write 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 an anonymous user. 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 allow 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.

12.3 Accessing data remotely with NFS - the command line view

To mount a remote filesystem 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 just 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 then you can use the mount command as follows:

```
mount -t nfs nfs_host:share_dir local_mount_dir
```

where:

`-t nfs`	Says to do the mount as an NFS mount. This is now optional because if you explicitly specify the directory to be mounted as host:directory the `mount` command knows that it is an NFS mount.
`nfs_host`	is the host that is exporting the filesystem to be shared.
`share_dir`	is the actual directory that is to be shared.
`local_mount_dir`	is the directory on the local host where the remote directory is going to be mounted. As mentioned earlier, this mount point must exist.

12.4 Allowing NFS access to data with GUI

In Caldera OpenLinux you can also mount the NFS resources using a graphical user interface. We explained this in 3.10.1, “Mounting an NFS volume” on page 110.

Chapter 13. 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!

13.1 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.

13.2 What can you do with Linux packet filtering?

With Linux packet filtering you can do many things. We 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 the 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 to this job.
- You can share your dial-up Internet connection with others.

13.3 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, enable **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 Caldera OpenLinux 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.

13.4 Network configuration for a packet filtering implementation

In this section we describe our lab network setup for implementing a packet filtering solution.

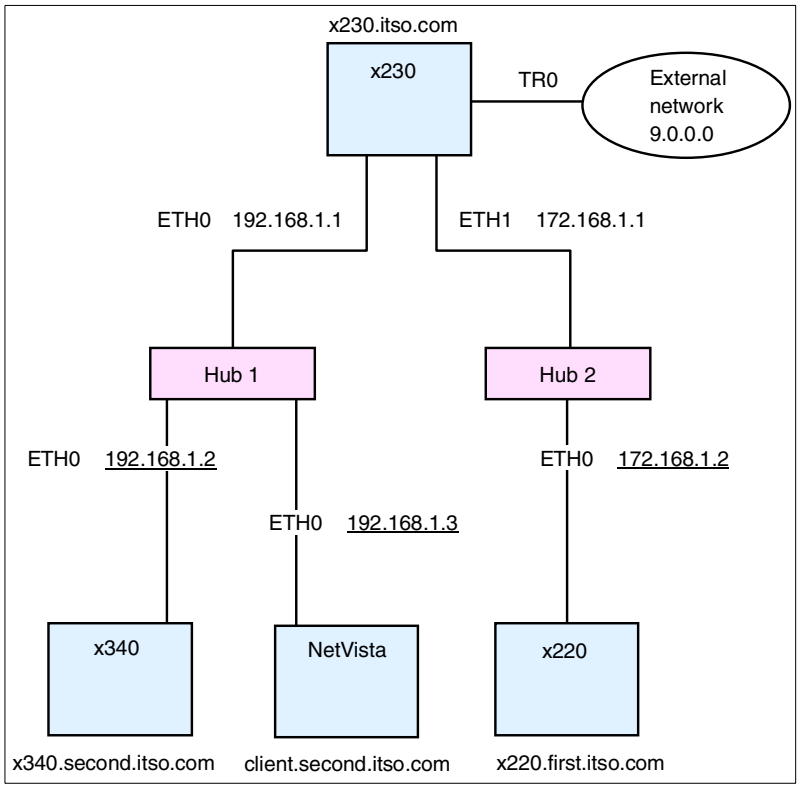

Figure 205. Lab network setup for firewall solution

Figure 205 shows our network setup:

- A x230 eServer 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
- A x340 eServer with one NIC and the following settings:
  - Eth0 - 192.168.1.2, default gateway 192.168.1.1
- A x220 eServer with one NIC and the following settings:
  - Eth0 - 172.168.1.2, default gateway 172.168.1.1
- A 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.

13.5 How to permanently enable IP Forwarding

In Caldera OpenLinux the network process is started by executing this script during the server startup:

```
/etc/rc.d/init/network
```

The IP Forwarding is not enabled by default. You can enable it by using the Webmin tool. To start the Webmin tool follow the instructions in 3.22, “System administration using Webmin” on page 129. After starting Webmin follow these steps to enable IP Forwarding:

1. In the main Webmin window select the **Hardware** tab, and you will see a window similar to Figure 206.

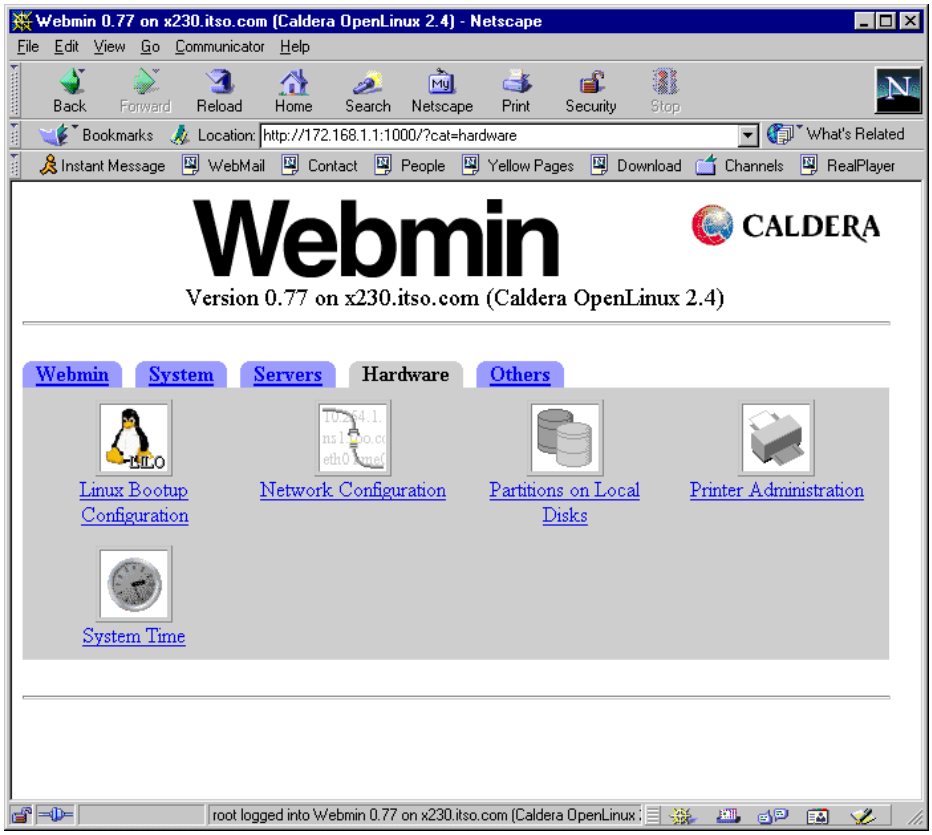

Figure 206. Webmin Hardware setup

2. Select **Network Configuration** and you will see a window similar to Figure 207.

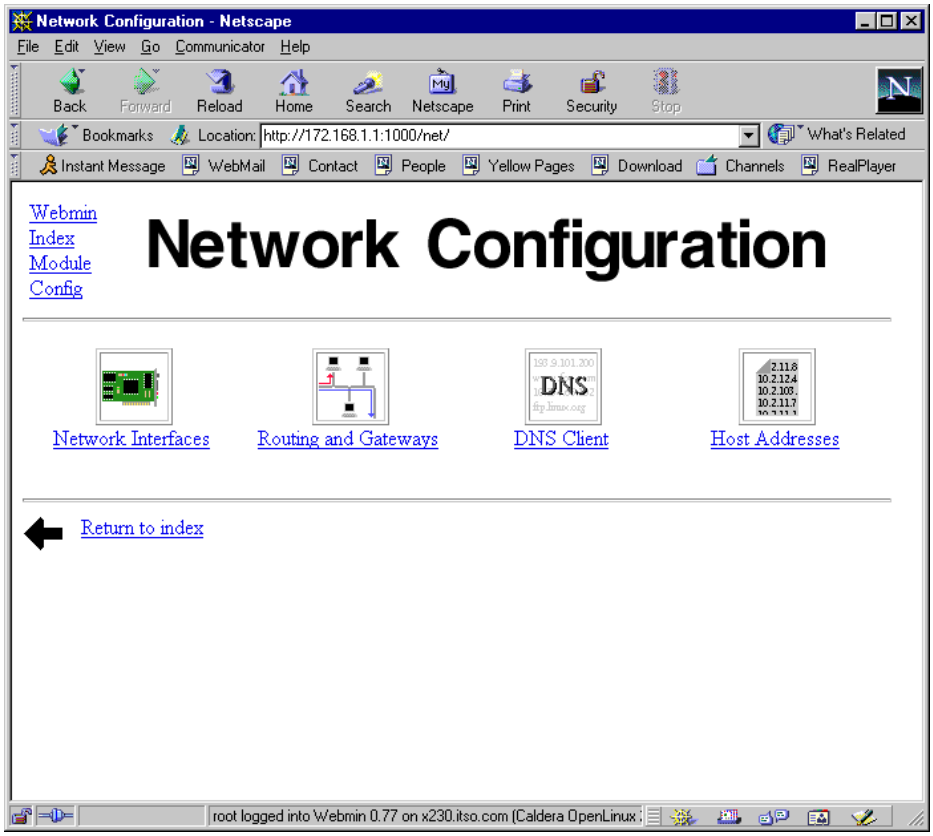

Figure 207. Webmin Network Configuration

3. Select **Routing and Gateways** and you will see a window similar to Figure 208.

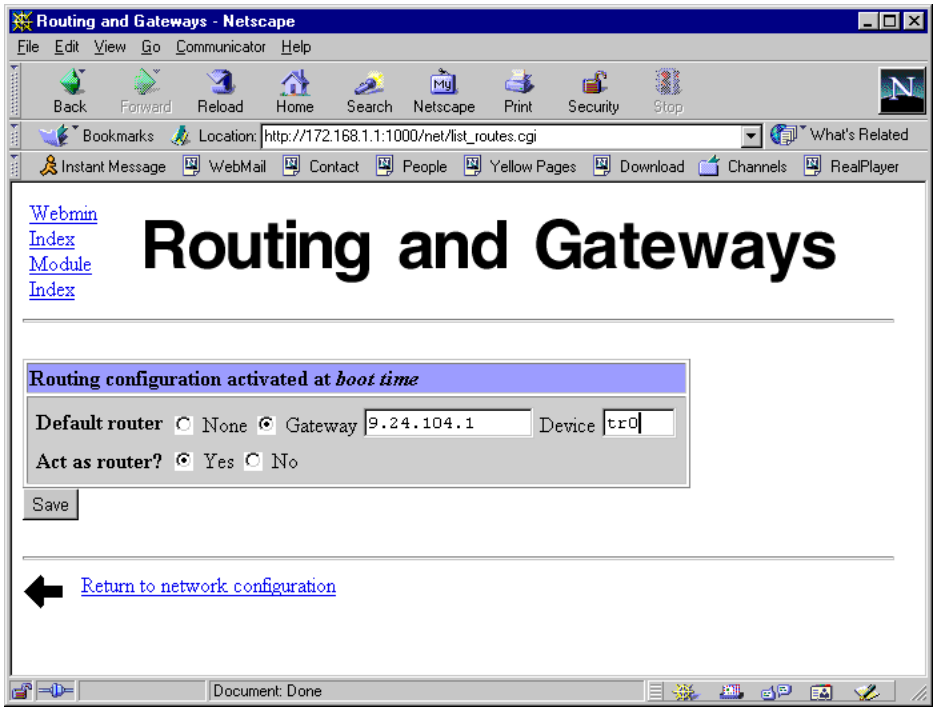

Figure 208. Webmin Routing and Gateways

4. To enable IP Forwarding select **Yes** at the **Act as router?**

After enabling IP Forwarding you need to restart network to activate the change. You can do this by executing the commands:

```
/etc/rc.d/init/network stop
/etc/rc.d/init/network start
```

You can also enable IP Forwarding by inserting the following line:

```
IPFORWARDING=yes
```

into the /etc/sysconfig/network file. After that you need to restart the network to activate the change.

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 209.

```
[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 209. PING after enabling IP Forwarding

13.6 Your first IP Chains success

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. How can we 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:

   ```
   /sbin/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:

   a. `8000` - timeout value for TCP sessions in seconds

   b. `20` - timeout value for TCP sessions after a FIN packet in seconds

   c. `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 210.

```
[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 210. FTP after IP Masquerading setup

You have just enabled access from internal networks to an external network.

13.7 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 211.

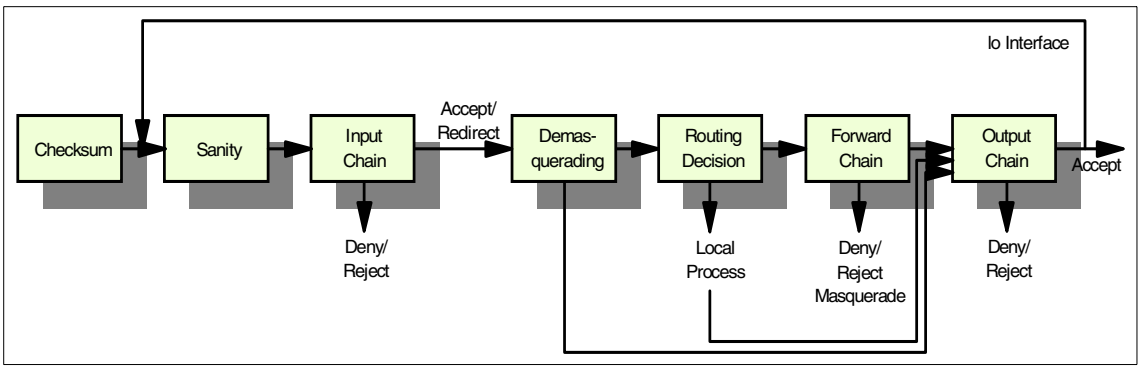

Figure 211. How the packet is traveling

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 a local process, they will go through the output chain with the interface set to “lo”. Then they will return to the input chain with the 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 211, 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 output similar to Figure 212.

```
[root@client /root]# ipchains -L
Chain input (policy ACCEPT):
Chain forward (policy ACCEPT):
Chain output (policy ACCEPT):
[root@client /root]#
```

Figure 212. Listing the default IP Chains policies

13.8 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 18.

Table 18. 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

To manipulate rules inside the chain, you can use the parameters explained in Table 19.

Table 19. 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 20.

Table 20. Parameters for managing masquerading

Parameter	Description
-M -L	List the currently masqueraded connections
-M -S	Set masquerading timeout values

13.8.1 How to create a rule

The most common syntax for the creating a new rule is:

```
/sbin/ipchains -A input -s source -p protocol -j action
```

The parameters are described in Table 21.

Table 21. IPChains parameters

Parameter	Description
-A	Append a new rule to the chain
source	IP address or hostname 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.

13.8.1.1 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 are gone.

13.8.2 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
```

13.9 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 14. 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.

14.1 Installing SSH

The Caldera OpenLinux eServer 2.3 does not include SSH by default because the US Government will not allow the import and export of strong encryption software.

Before installing SSH on Caldera OpenLinux eServer 2.3 you need to download the following source code:

```
OpenSSL - http://www.openssl.org/source/openssl-0.9.6.tar.gz
OpenSSH -
ftp://ftp1.usa.openbsd.org/pub/OpenBSD/OpenSSH/portable/openssh-2.3.0p1.ta
r.gz
```

14.1.1 Installing OpenSSL

Now change to the directory where you downloaded the source code of OpenSSL and follow the steps below:

1. Untar the source with the command:

   ```
   tar -zxvf openssl-0.9.6.tar.gz
   ```

2. Change to the resulting directory and compile the package with the commands:

   ```
   cd openssl-0.9.6
   ./config --prefix=/usr
   make
   make install
   ```

14.1.2 Installing OpenSSH

Change to the directory where you downloaded the source code of OpenSSH and follow the steps below:

1. Untar the source with the command:

```
tar -zxvf openssh-2.3.0p1.tar.gz
```

2. Change to the resulting directory and compile the package with the commands:

```
cd openssl-0.9.6
./config --prefix=/usr --sysconfdir=/etc/ssh
make
make install
```

3. Copy the sshd.pam.generic file to the /etc/pam.d directory with the commands:

```
cd contrib
cp sshd.pam.generic /etc/pam.d/sshd
```

14.2 Configuring SSH

Before you can start configuring you need to change the sshd file in /etc/pam.d so that you replace all instances of `pam_unix.so` with `pam_pwdb.so`. After the change, the sshd.pam.generic file should look similar to the following:

```
#%PAM-1.0
#auth       required     /lib/security/pam_unix.so shadow nodelay
auth       required     /lib/security/pam_pwdb.so shadow nodelay
auth       required     /lib/security/pam_nologin.so
#account    required     /lib/security/pam_unix.so
account    required     /lib/security/pam_pwdb.so
password   required     /lib/security/pam_cracklib.so
#password   required     /lib/security/pam_unix.so shadow nullok
use_authtok
password   required     /lib/security/pam_pwdb.so shadow nullok
use_authtok
#session    required     /lib/security/pam_unix.so
session    required     /lib/security/pam_pwdb.so
session    required     /lib/security/pam_limits.so
```

To set up a secure system, we need to tell the SSH daemon not to allow a fallback to password authentication, since 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 logging 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 adaptors to the ListenOn parameter. Using an IP address of 0.0.0.0 will tell the SSH daemon to listen on all network devices in the system.

14.2.1 Host key generation

After the installation of OpenSSH, it will automatically create a random host key to uniquely identify the system.

```
Generating RSA keys:  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:
85:72:08:d4:91:6b:e5:69:9f:2f:ad:d0:c2:1e:a7:c5 root@x220.first.itso.com
Generating DSA parameter and key.
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:
56:19:79:71:ec:25:fb:19:75:2e:1f:1c:ba:b1:27:6a root@x220.first.itso.com
```

Figure 213. After installing the OpenSSH

14.2.2 SSHD server dameon

To run the sshd server daemon type `sshd` at a command line and press Enter. To have the sshd server start automatically upon boot up, do the following:

1. Edit the /etc/rc.d/rc.local file and add `sshd` to the bottom of the file.
2. Save the changes and reboot.

14.2.3 User key generation

Once this has been done we need to create a public/private key pair for users. To do this a user must run `ssh-keygen` at a command prompt. You will be asked where to store your key pair. The default values are fine for this.

```
[root@x220 pam.d]# ssh-keygen
Generating RSA keys:  Key generation complete.
Enter file in which to save the key (/root/.ssh/identity):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/identity.
Your public key has been saved in /root/.ssh/identity.pub.
The key fingerprint is:
5e:20:71:9c:c6:dd:e9:45:9f:9a:a1:1b:9a:53:cf:7f root@x220.first.itso.com
```

Figure 214. 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.

14.2.4 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_hosts file. This way you can keep on adding public keys if you have more than one combination.

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!

On another host, copy your private key to ~/.ssh/identity. Make sure no one can read this file except you by setting the correct permissions:

```
chmod 400 ~/.ssh/identity
```

This will make sure only the owner of the file can read 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 215. 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 15. SNMP

In this section we will be taking a brief look at SNMP. For our examples in Linux we will be using packages associated with the net-snmp project, formerly known as UCD-snmp. This project was originally based on the Carnegie Mellon University and University of California at Davis SNMP implementations. The project has grown and changed significantly since its inception. Additional documentation and source code is available at:

```
http://net-snmp.sourceforge.net/
```

15.1 SNMP - What is it?

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 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.

15.2 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.

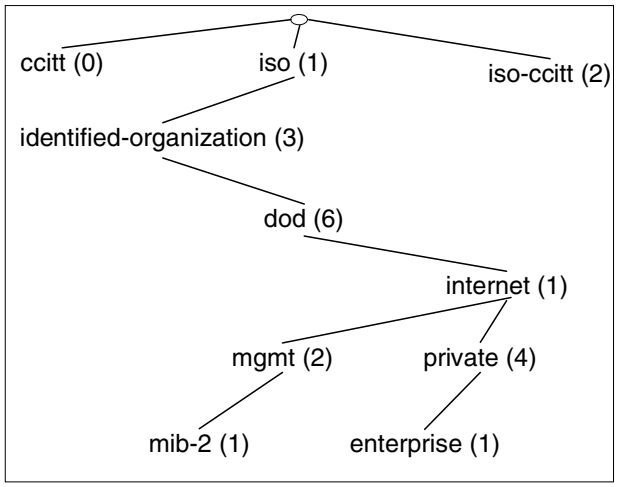

Figure 216. Partial MIB tree illustration

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 (for example, “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 subset.

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.

15.3 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. 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 at:

`http://www.hylafax.org`

15.4 Implementation on Linux

The Caldera OpenLinux 2.3 eServer distribution includes the following RPMs:

ucd-snmp-4.0.1-3

ucd-snmp-devel-4.0.1-3

ucd-snmp-utils-4.0.1-3

If these packages are not installed, mount the Caldera OpenLinux 23. eServer CDROM and type the following command:

```
rpm -ivh /mnt/cdrom/col/install/RPMS/ucd*
```

After installing the packages we need to edit or create the /etc/snmp/snmpd.conf file.

Our sample snmpd.conf illustrates a very limited implementation. We grant read access to anyone on the local network (192.168.1.0) using the community string of public.

```
############################################################
# snmpd.conf:
############################################################
com2sec mynetwork 192.168.1.0/24 public
group MyROGroup  any        mynetwork
```

```
view all    included  .1                                80
access MyROGroup ""      any       noauth    0      all    none   none
syslocation Seattle
syscontact jhaskins@ibm.com
#############################################################
```

What does it mean:

`com2sec NAME SOURCE COMMUNITY`

This entry specifies the mapping from a source/community pair to a security name. `SOURCE` can be a hostname, a subnet, or the word "default". A subnet can be specified as IP/MASK or IP/BITS. The first source/community combination that matches the incoming packet is selected.

In our example we used:

```
com2sec mynetwork 192.168.1.0/24 public
```

This defines mynetwork as 192.168.1.0/24 and maps it to the public community string. Although public is the default community string for read-only access, we recommend changing the community string to something other than public.

`group NAME MODEL SECURITY`

This entry defines the mapping from security-model/securityname to a group. `MODEL` is one of v1, v2c, or usm.

In our example we used:

```
group MyROGroup  any        mynetwork
```

This entry defines the group MyROGroup with any of the available security models and maps it to mynetwork.

`view NAME TYPE SUBTREE [MASK]`

This entry defines the named view. `TYPE` is either included or excluded. `MASK` is a list of hex octets, separated by '.' or ':'. The mask defaults to `ff` if not specified.

In our example we used:

```
view all    included  .1
```

This entry defines the view as all and includes everything (.1).

`access NAME CONTEXT MODEL LEVEL PREFX READ WRITE NOTIFY`

This entry maps the group/security model/security level to a view. `MODEL` is one of any, v1, v2c, or usm. `LEVEL` is one of noauth, auth, or priv. `PREFX`

specifies how context should be match against the CONTEXT of the incoming pdu, either exact or prefix. READ, WRITE, and NOTIFY specifies the view to be used for the corresponding access.

In our example we used:

```
access MyROGroup ""      any       noauth    0       all     none    none
```

This entry defines the access for MyROGroup with any security model, the ability to read all, but no access to write or notify.

We also define a location and a contact for the system. This agent is responsible for:

```
syslocation Seattle, WA USA
syscontact jhaskins@ibm.com
```

After creating our basic snmpd.conf file, start snmpd:

```
/etc/init.d/snmpd 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 = Linux m10A 2.2.16-22 #1 SMP Mon Oct 30 14:36:08 EST 2000 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 net-snmp(usc-snmp) packages 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 up traps to be sent to a specified host. The documentation included with net-snmp is quite thorough. Take a look.

15.4.1 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.

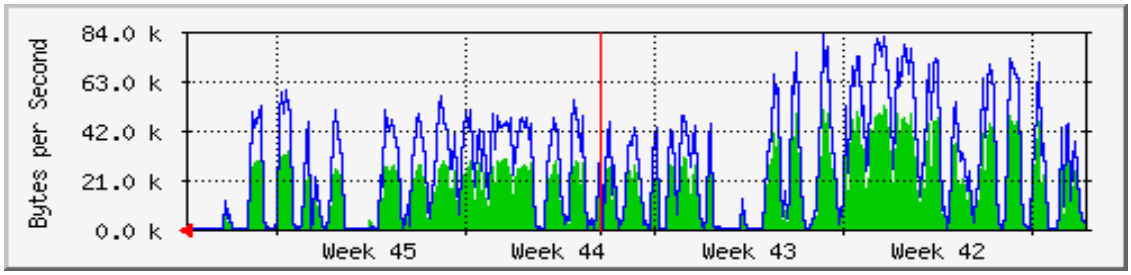

Figure 217. MRTG in action

MRTG requires Perl to be installed. If the Perl package is not installed, simply mount the distribution media and install the following RPM:

perl-5.6.0-9

Although MRTG can be used without a Web server, we recommend you install the Apache Web server. See Chapter 8, “Apache and IBM HTTP Servers” on page 199.

Installation of MRTG on Caldera OpenLinux 2.3 eServer is simple. Download the latest tarball from `http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/pub/` to the /tmp directory.

For our example we are using the mrtg.2.9.4.tar.gz.

Unpack the gzipped tarball, run the configure script, and then make the package:

```
cd /tmp
tar -xvzf mrtg.2.9.4.tar.gz
cd mrtg.2.9.4
```

```
./configure
make
make install
```

This will install MRTG in /usr/local/mrtg-2/

Now 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/mrtg-2/bin
./cfgmaker public@localhost >mrtg.cnf
```

You should see the following output:

```
--base: Get Device Info on public@localhost
--base: Vendor Id:
--base: Populating confcache
--base: Get Interface Info
--base: Walking ifIndex
--base: Walking ifType
--base: Walking ifSpeed
--base: Walking ifAdminStatus
--base: Walking ifOperStatus
```

We need to edit the newly created mrtg.cnf. Under the section Global config options, uncomment and change the WorkDir entry for UNIX to the default HTML directory with the subdirectory mrtg. (/var/www/html/mrtg). In our example we are using the Apache Web server, which serves Web pages from /var/www/html. Save the file.

Sample mrtg.cnf file:

```
##################################################################
# System: phu
# Description: Linux phu 2.2.16-22smp #1 SMP Tue Aug 22 16:39:21 EDT 2000
# 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 ./mrtg.cnf
Rateup WARNING: /usr/local/mrtg-2/bin/rateup could not read the primary log
file for localhost_4
Rateup WARNING: /usr/local/mrtg-2/bin/rateup The backup log file for
localhost_4 was invalid as well
Rateup WARNING: /usr/local/mrtg-2/bin/rateup Can't remove localhost_4.old
updating log file
Rateup WARNING: /usr/local/mrtg-2/bin/rateup Can't rename localhost_4.log
to localhost_4.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 /var/www/html/mrtg

localhost_1-day.png
localhost_1-month.png
localhost_1-week.png
localhost_1-year.png
localhost_1.html
localhost_1.log
localhost_1.old
```

Make sure your Web server is running and point your Web browser to `http://localhost/mrtg/localhost_1.html`.

You should see a Web page similar to Figure 218 on page 294 (with less data of course):

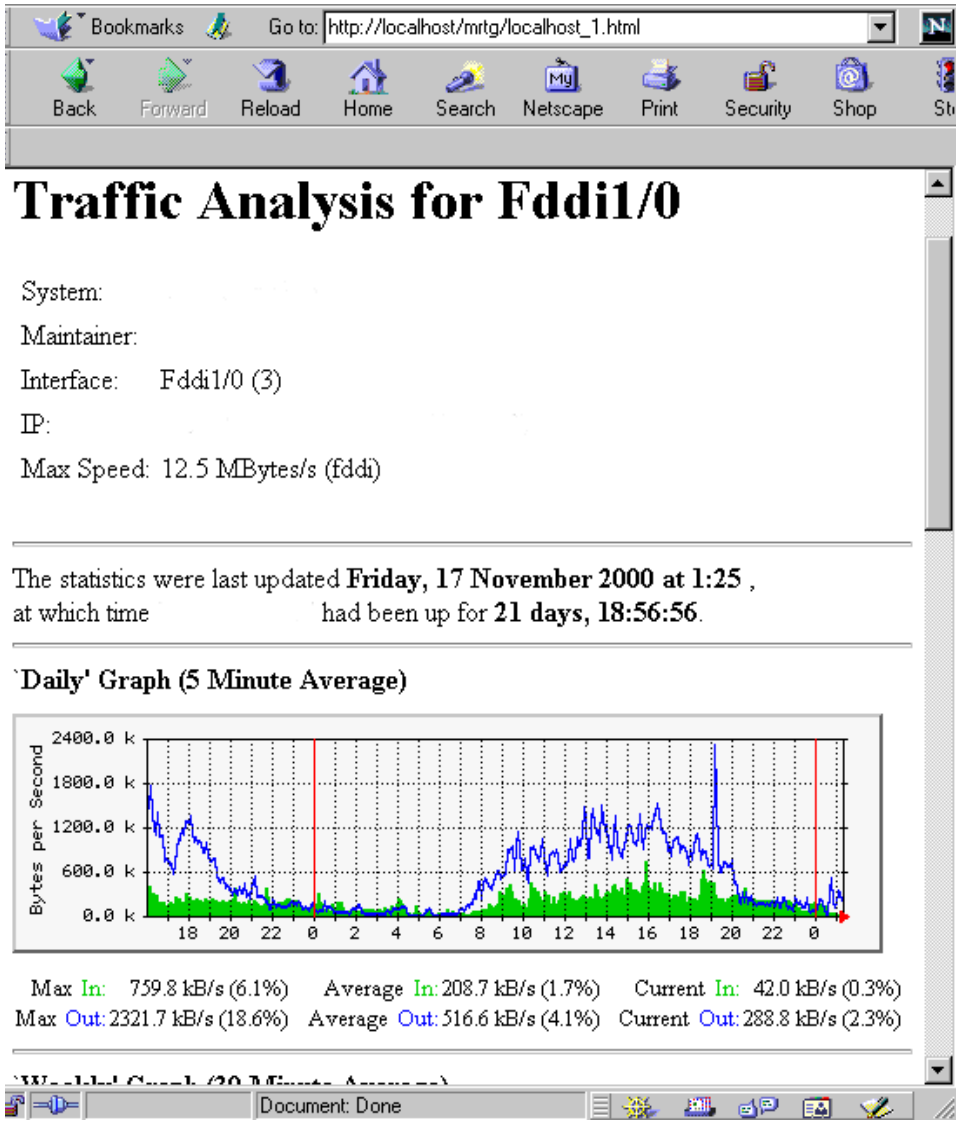

Figure 218. A sample MRTG page running for more than 24 hours on a FDDI interface

The final step is to automate the running of MRTG. Add the following entry to your crontab:

```
0/5 * * * * /usr/local/mrtg-2/bin/mrtg /usr/local/mrtg-2/bin/mrtg.cnf
```

Now MRTG will run every five minutes and update your Web page.

There is so much more you can do with this great package and SNMP. Check out the MRTG Web page for examples and documentation:

http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/mrtg.html

15.4.2 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

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 Caldera OpenLinux 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:

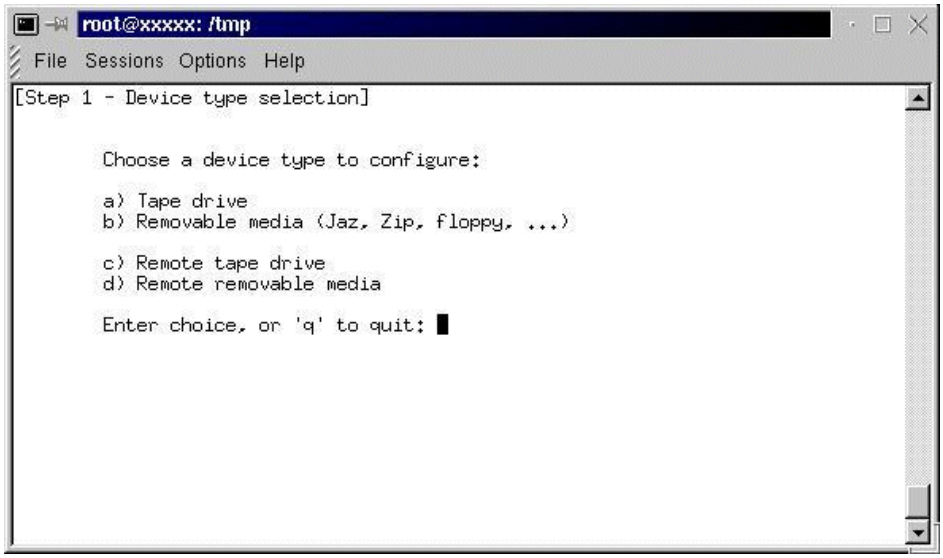

Figure 219. Selecting your backup devices

Enter the letter for your backup device and answer the following questions appropriate for your device.

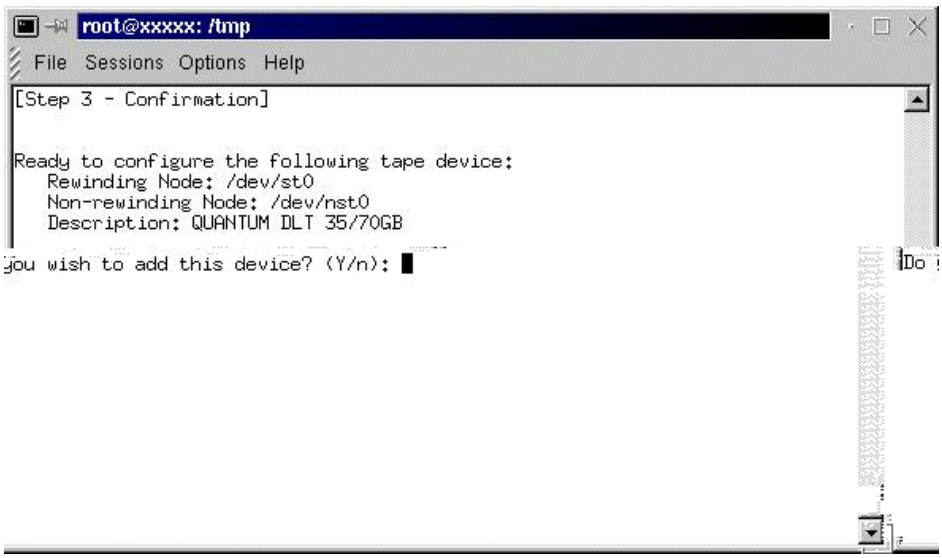

Figure 220. You have entered your backup devices

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:

```
xbru
```

Figure 221. XBRU window

You will see a window similar to Figure 221.

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 222).

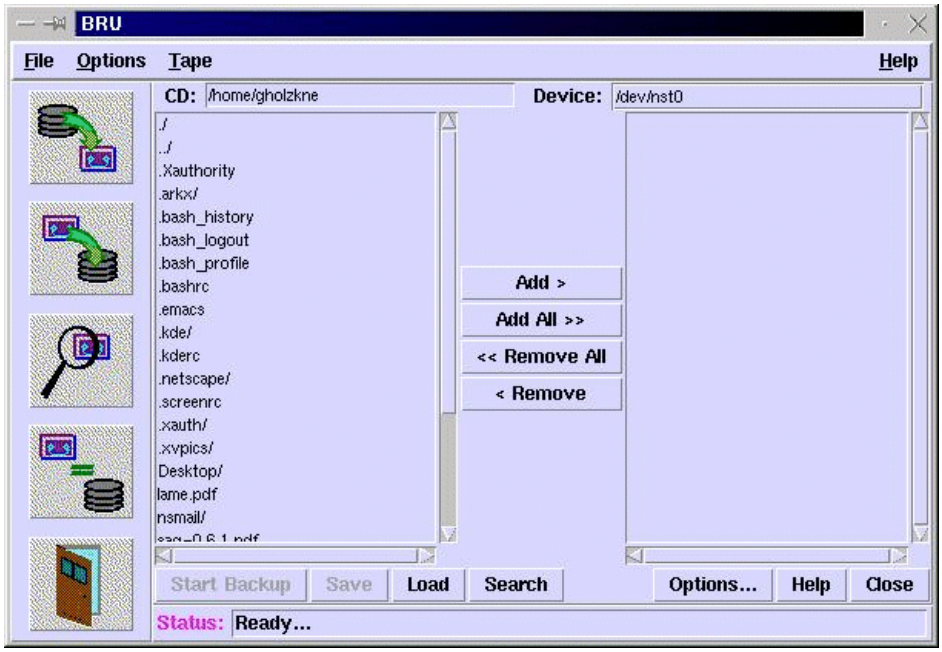

Figure 222. Creating an archive

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 223) 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.

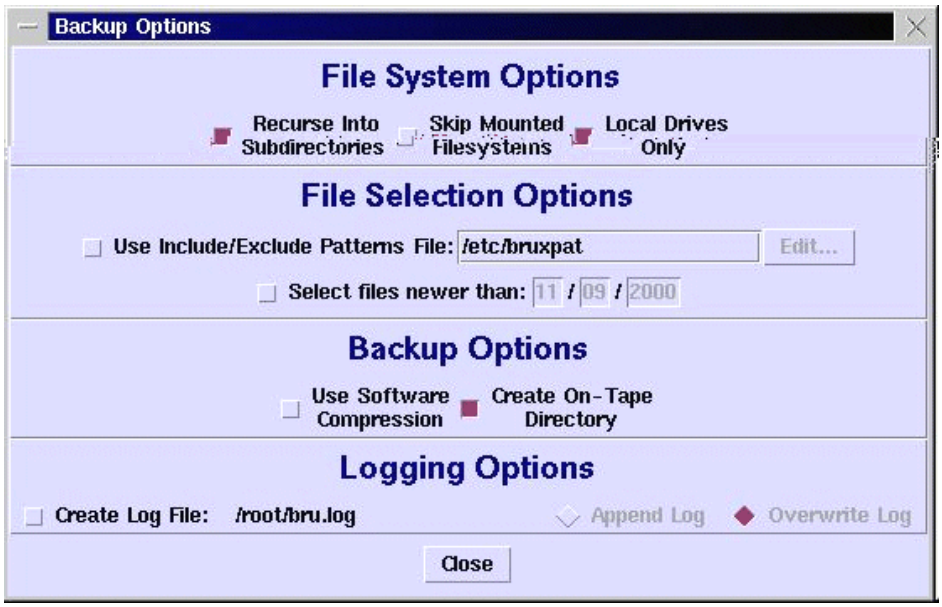

Figure 223. Dialog for backup options

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.

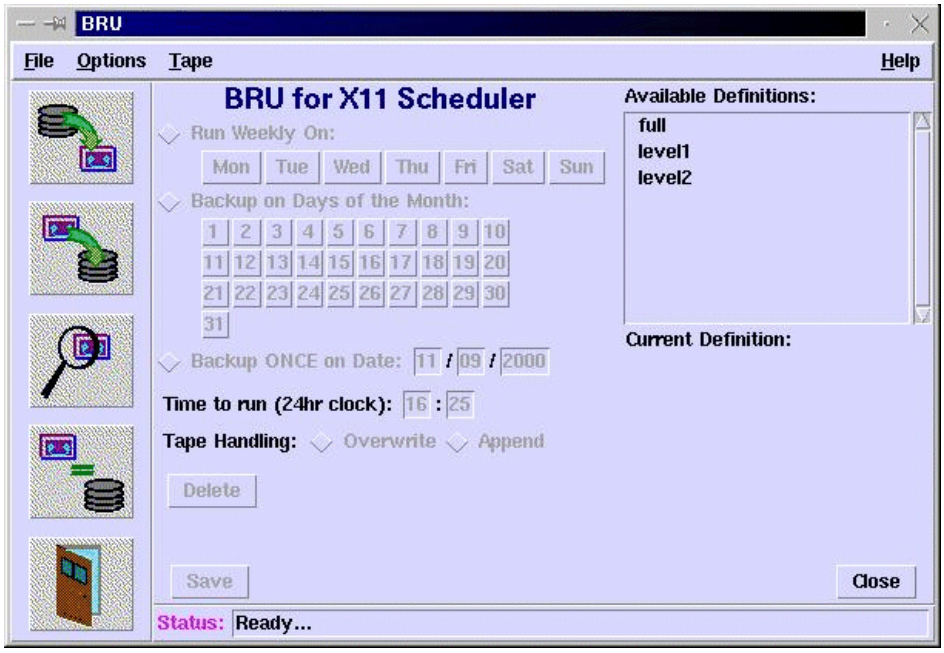

Figure 224. Scheduler

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 301. 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 222) 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 O**ptions** 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 225.

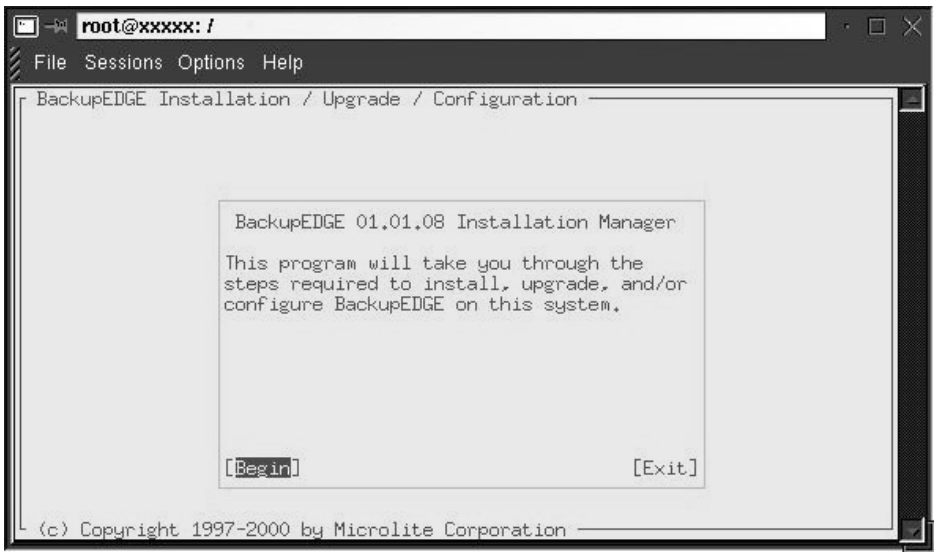

Figure 225. Start of installation dialog

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 edgemenu program by executing command:

   `edgemenu`

   You will see a window similar to Figure 226.

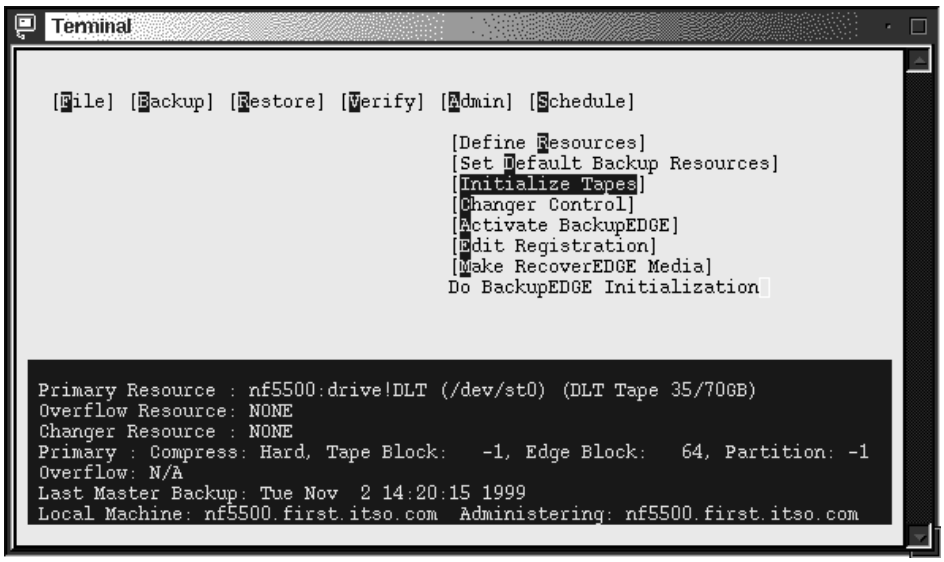

Figure 226. BackupEdge main menu

2. In the Admin menu select **Initialize Tapes.** You will see a window similar to Figure 227.

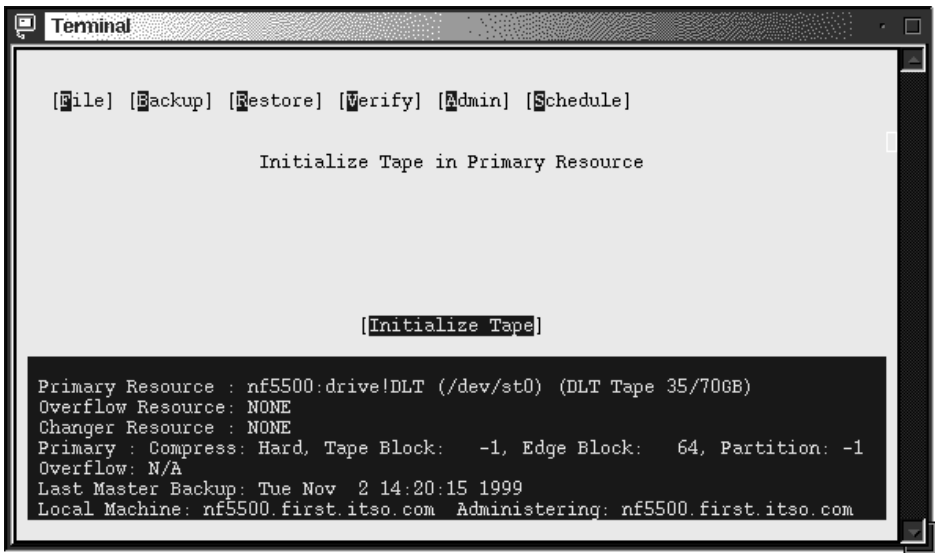

Figure 227. Initializing the tape

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 228.

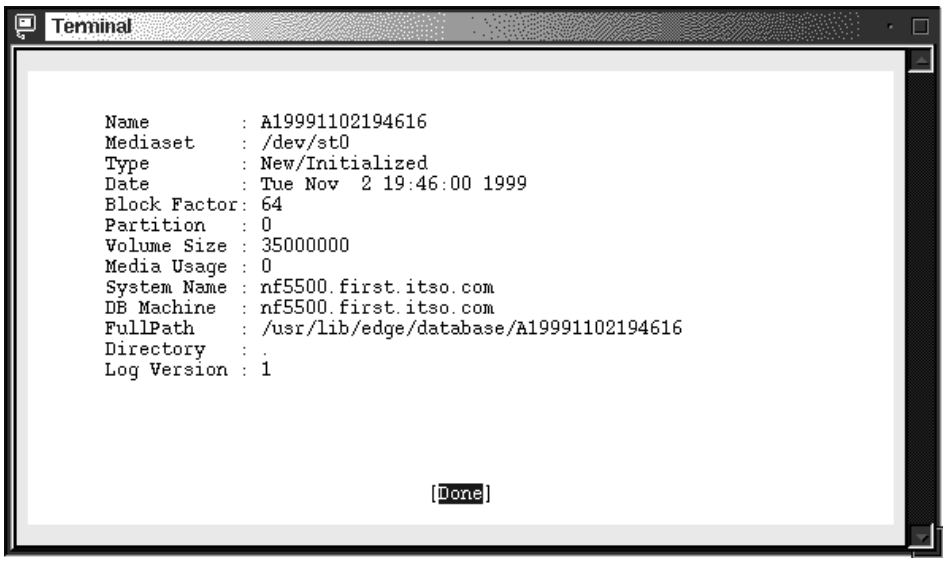

Figure 228. Tape information

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 229.

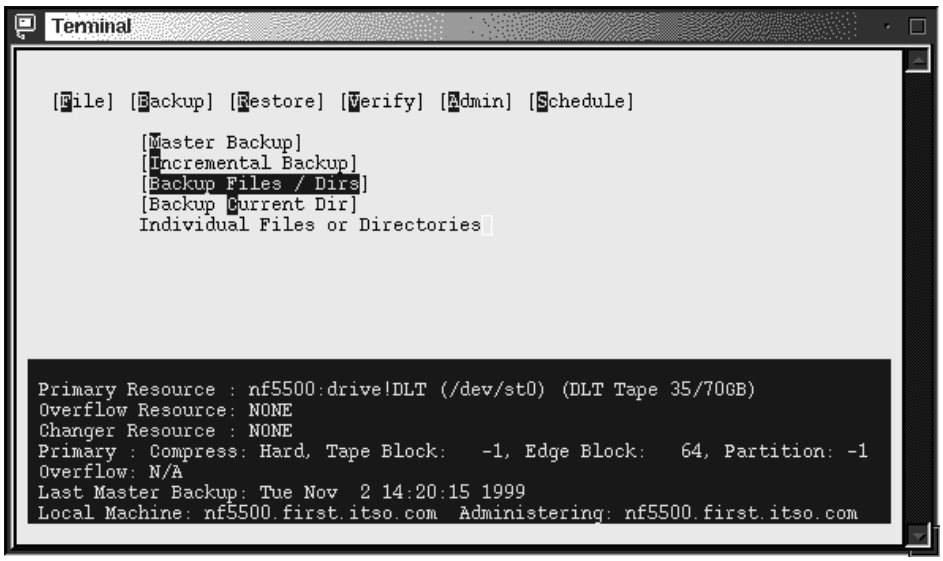

Figure 229. Starting the backup

2. In the Backup menu select **Backup Files / Dirs,** and you will see a window similar to Figure 230.

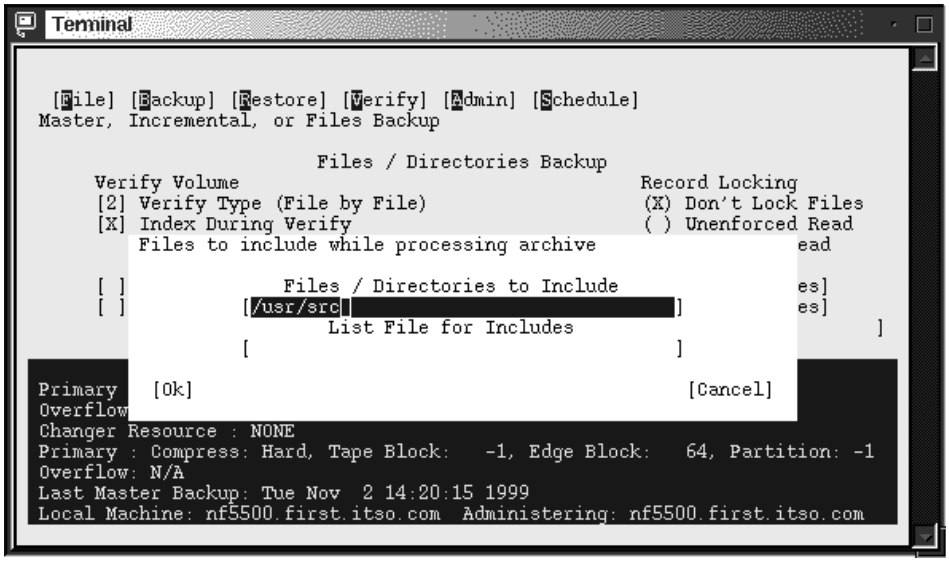

Figure 230. Selecting source for backup

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 231.

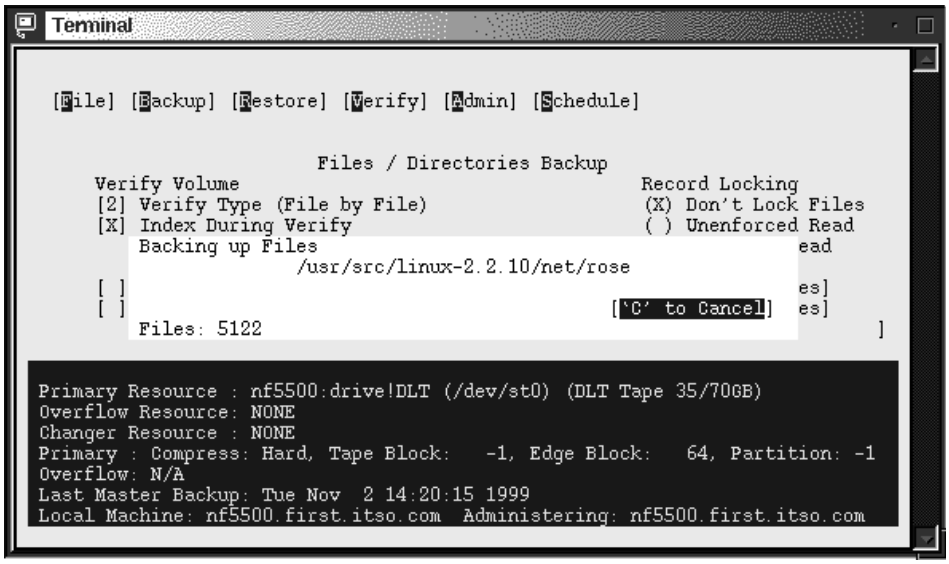

Figure 231. Backup in progress

After the backup is finished you will see a window similar to Figure 232.

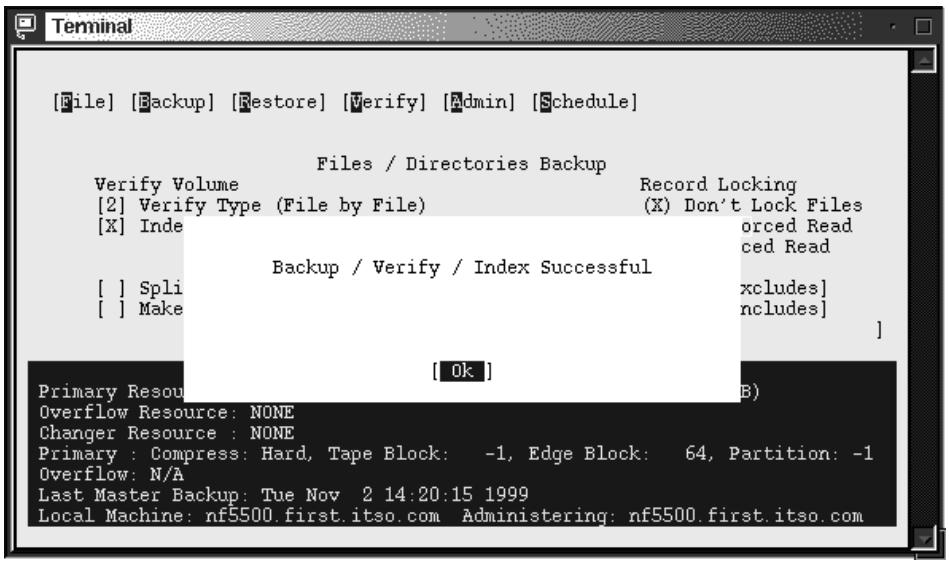

Figure 232. Backup completed

You will also see the backup report similar to Figure 233.

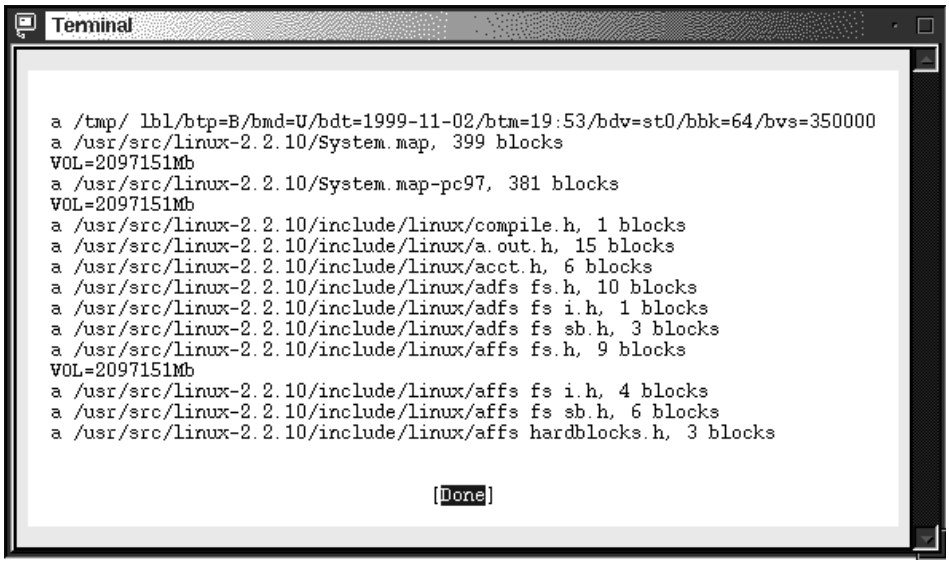

Figure 233. Backup report

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 229. Select **Restore** and a window similar to Figure 234.

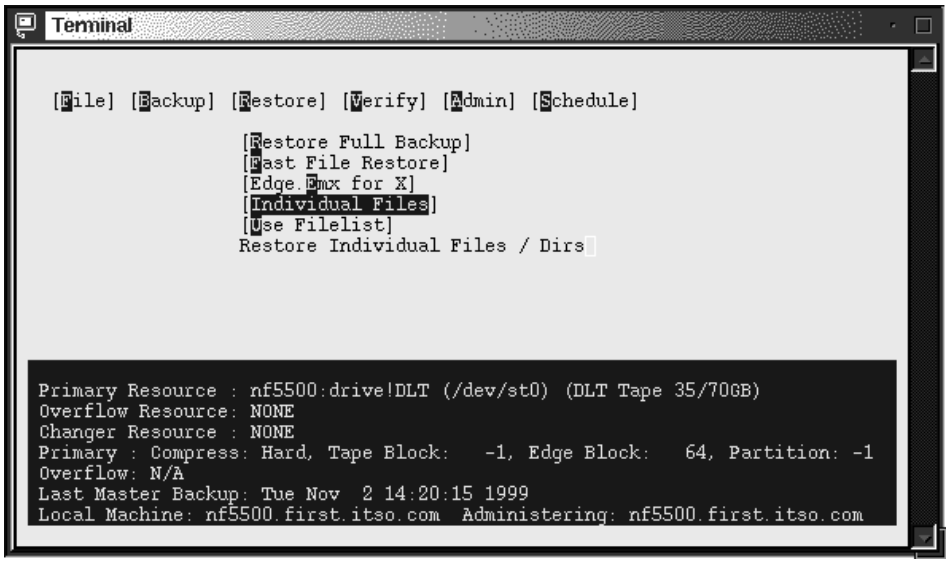

Figure 234. Starting the recovery

2. Select **Restore** > **Individual Files**, and you will see a window similar to Figure 230 on page 311.
3. Select the files or directories to restore. Select **OK** to continue, and you will see a window similar to Figure 235.

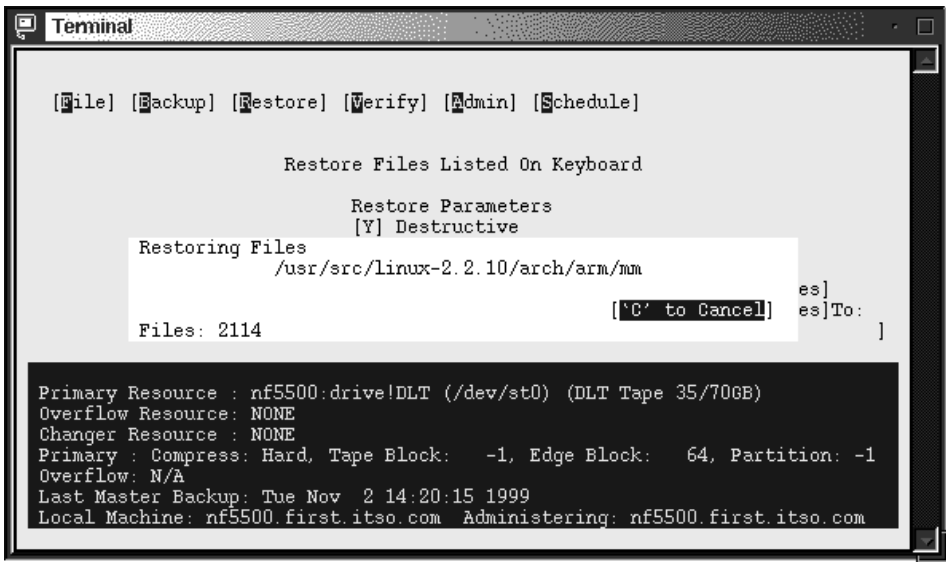

Figure 235. Recovery in progress

When the recovery is completed you will see a window similar to Figure 236.

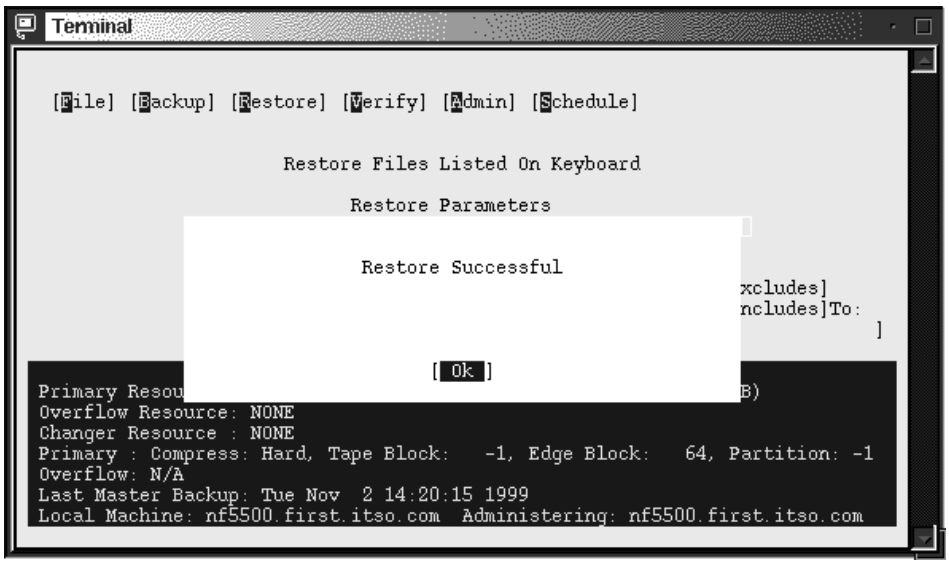

Figure 236. Recovery completed

Select **OK** to continue and you will see a recovery report similar to Figure 237.

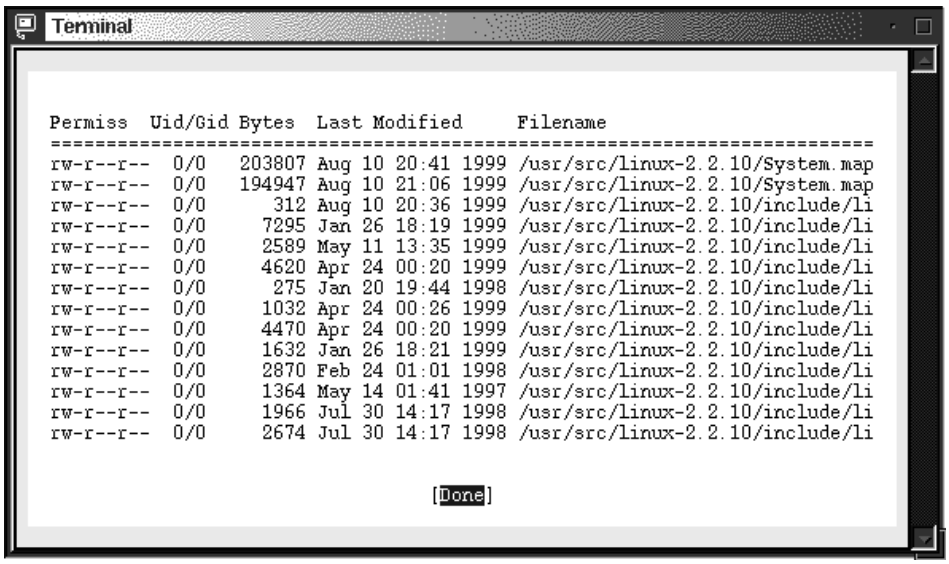

Figure 237. Recovery report

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 229 on page 311.

2. Select **Backup** > **Master Backup,** and you will see a window similar to Figure 238.

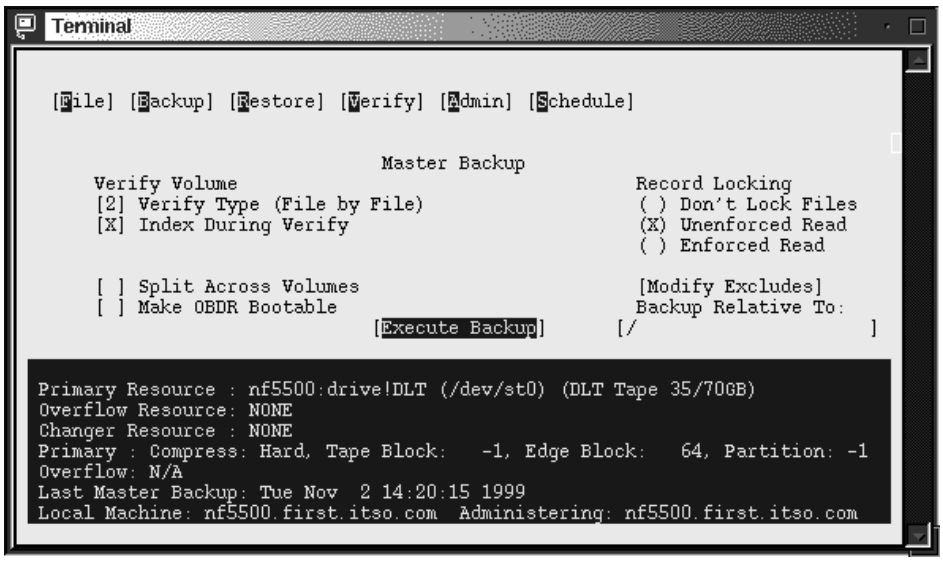

Figure 238. Starting the master backup

3. Choose the options you want and select **Execute Backup** to start the backup. You will see a window similar to Figure 239.

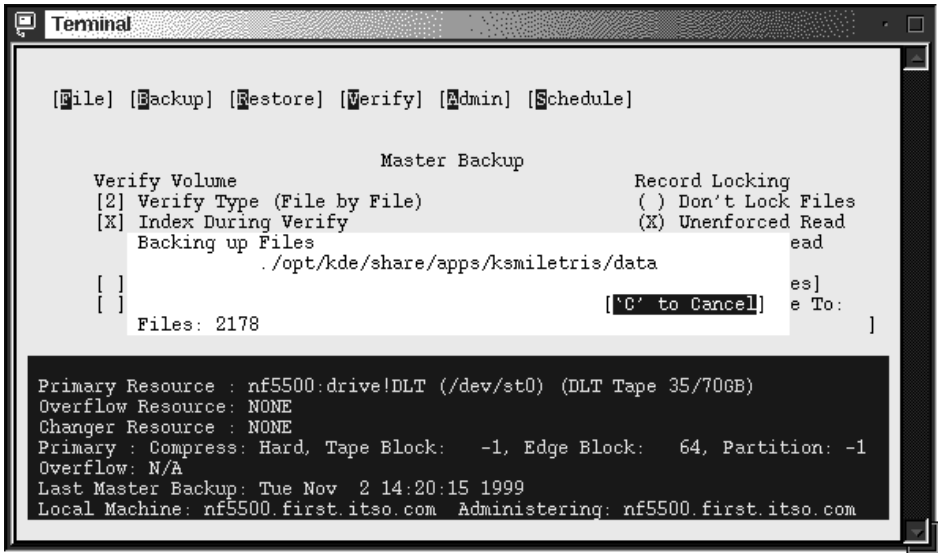

Figure 239. Master backup in progress

When the backup is finished you will see a window similar to Figure 240.

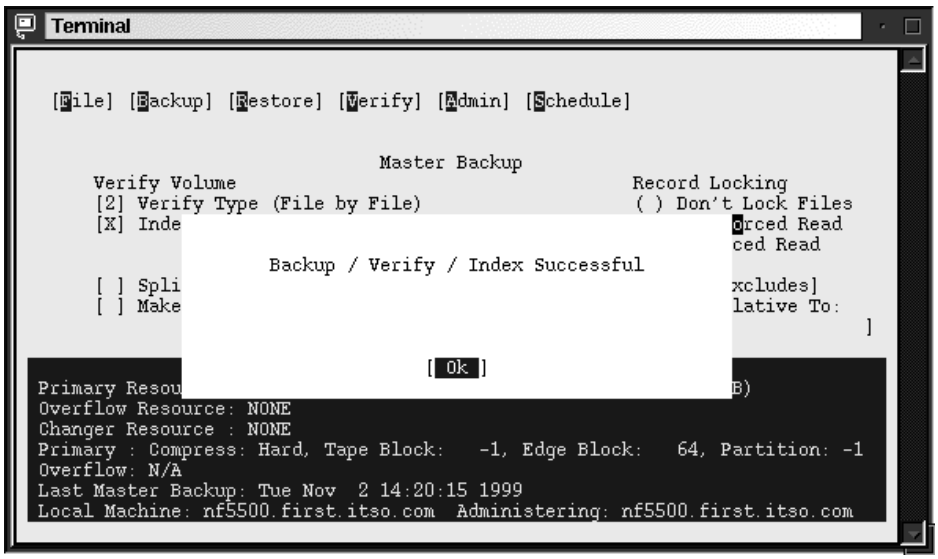

Figure 240. Master backup completed

Select **OK** to finish the operation, and you will see a backup report similar to Figure 241.

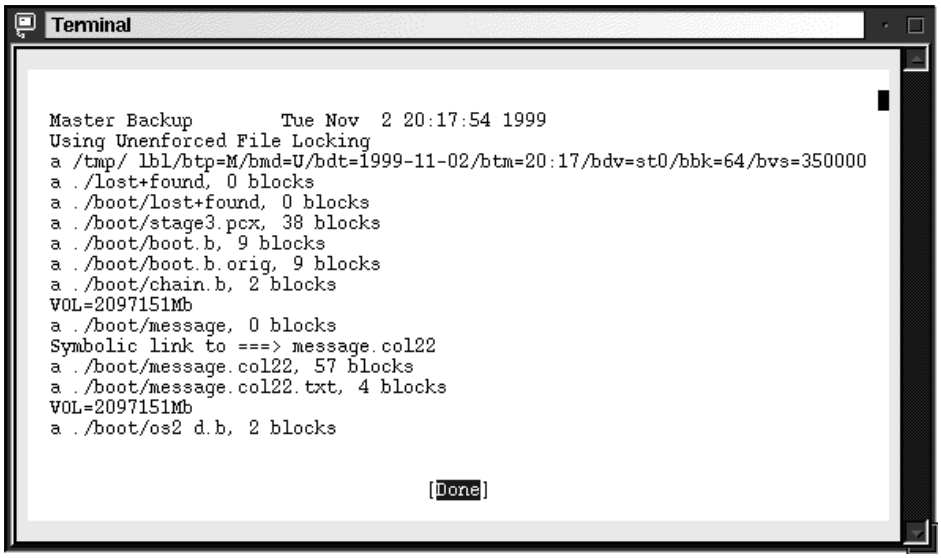

Figure 241. Master backup report

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 edgemenu utility. When you start the utility and choose **Restore** you will see a window similar to Figure 242.

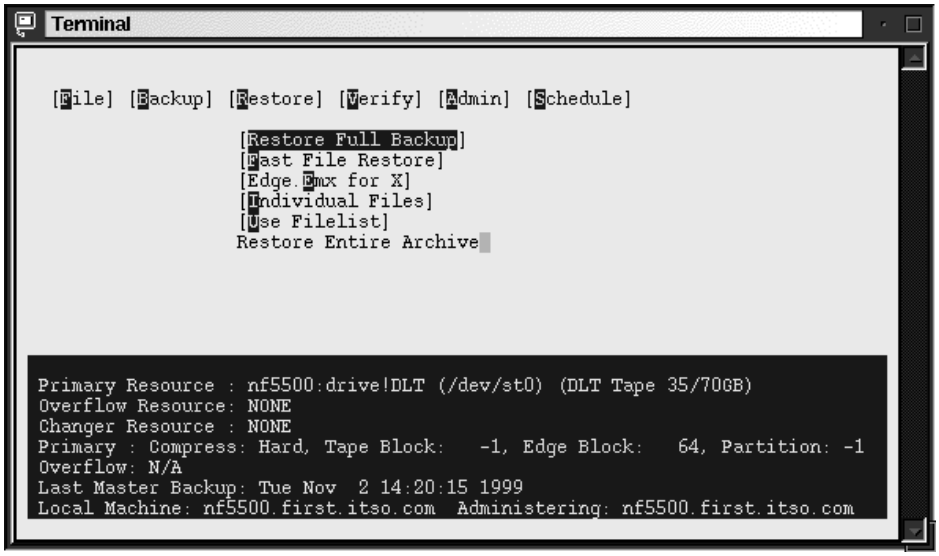

Figure 242. Starting restore full backup

Select **Restore > Restore Full Backup** and you will see a window similar to Figure 243.

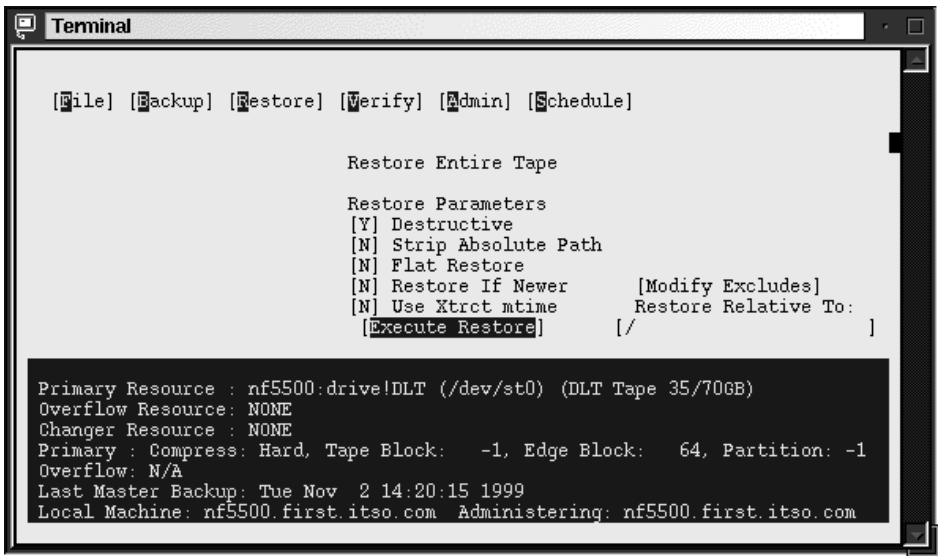

Figure 243. Full backup restore options

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 244.

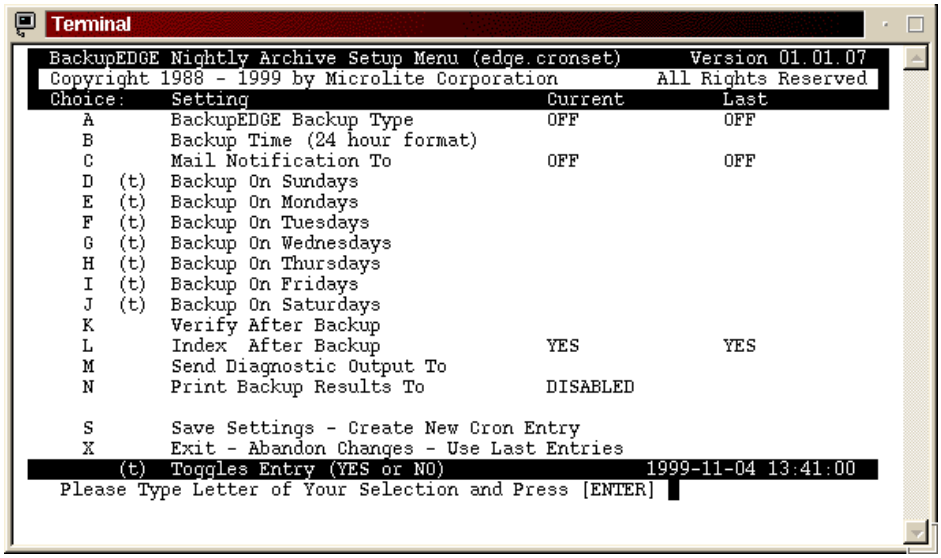

Figure 244. Schedule setup

3. 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 245.

Figure 245. Defining the type of backup

4. 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.

5. Next you need to specify the time of everyday backup by selecting **B** and pressing Enter. You will see a window similar to Figure 246.

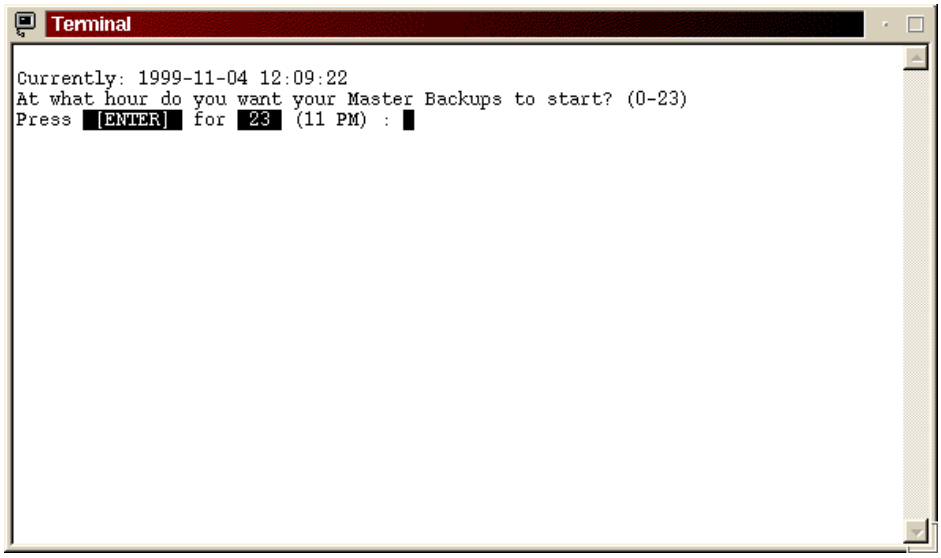

Figure 246. Setting the time

6. Define the time for your backups. You will see a window similar to Figure 247.

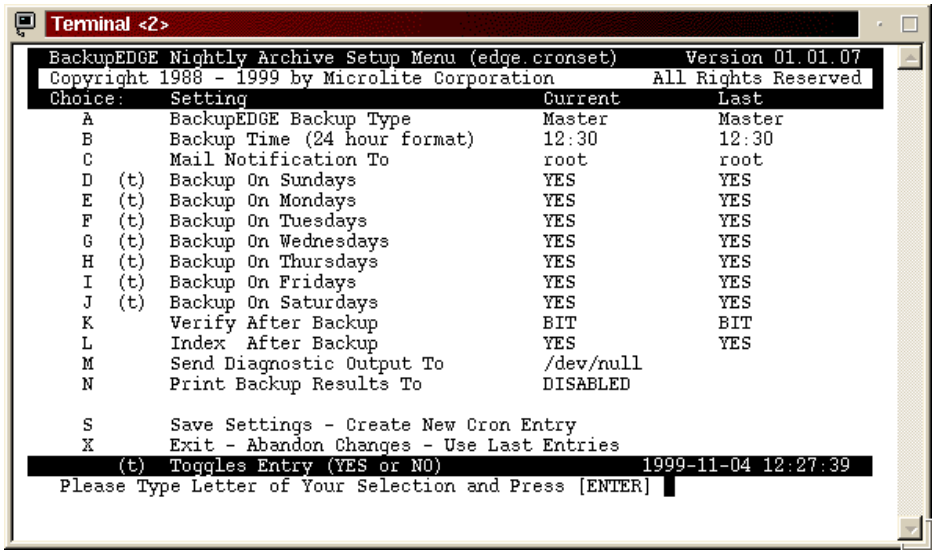

Figure 247. After schedule definition

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 248.

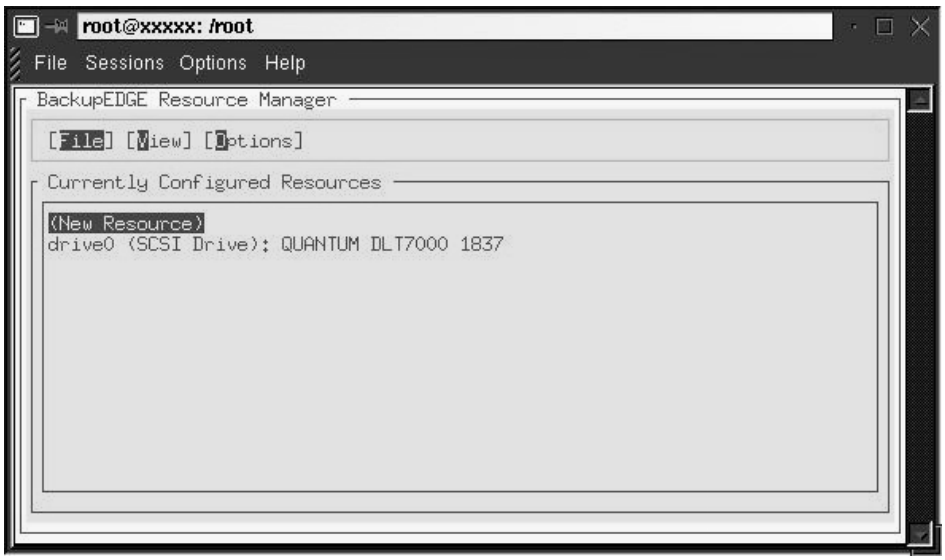

Figure 248. Starting the resource manager

2. Select **New Resource** and press Enter. You will see a window similar to Figure 249.

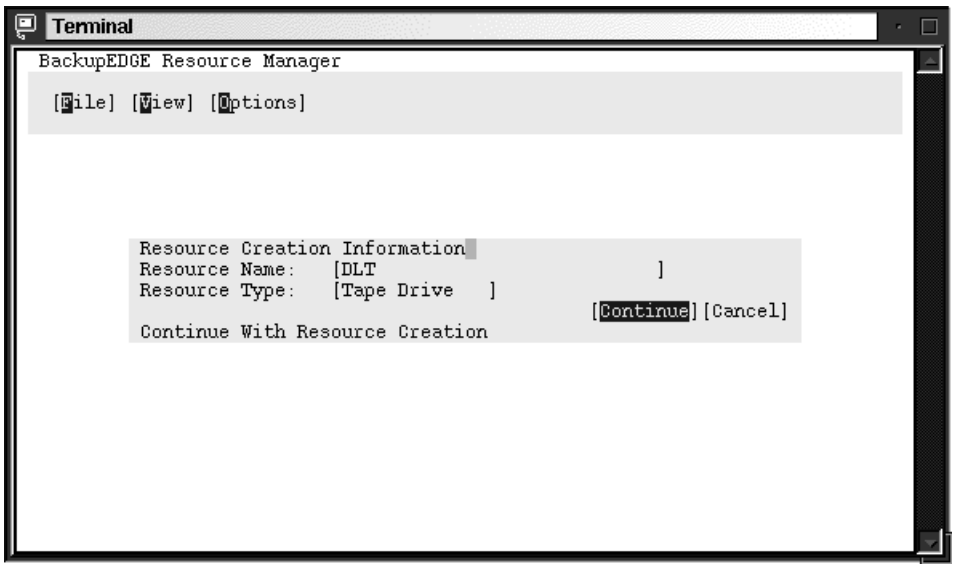

Figure 249. Defining the resource name

3. Type in the resource name and select a resource type. Select **Continue** to go on. You will see a window similar to Figure 250.

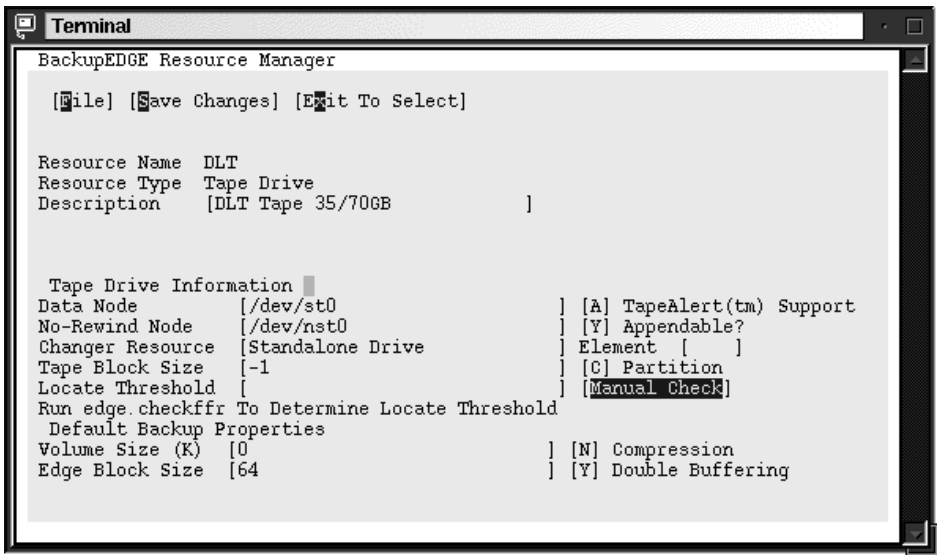

Figure 250. Parameters for the tape

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 251.

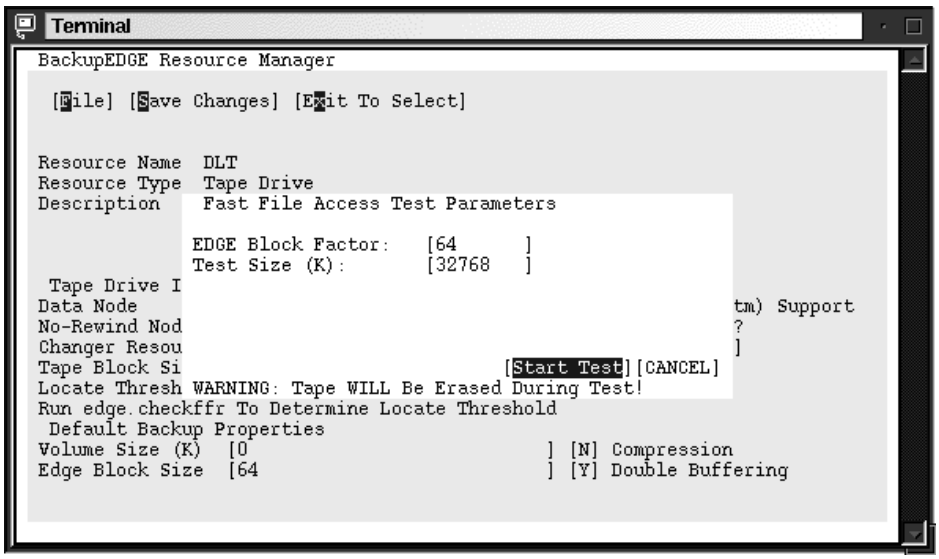

Figure 251. Setting the parameters for tests

6. Here you can select the block factor and the test size. Select **Start Test** to continue. You will see a window similar to Figure 252.

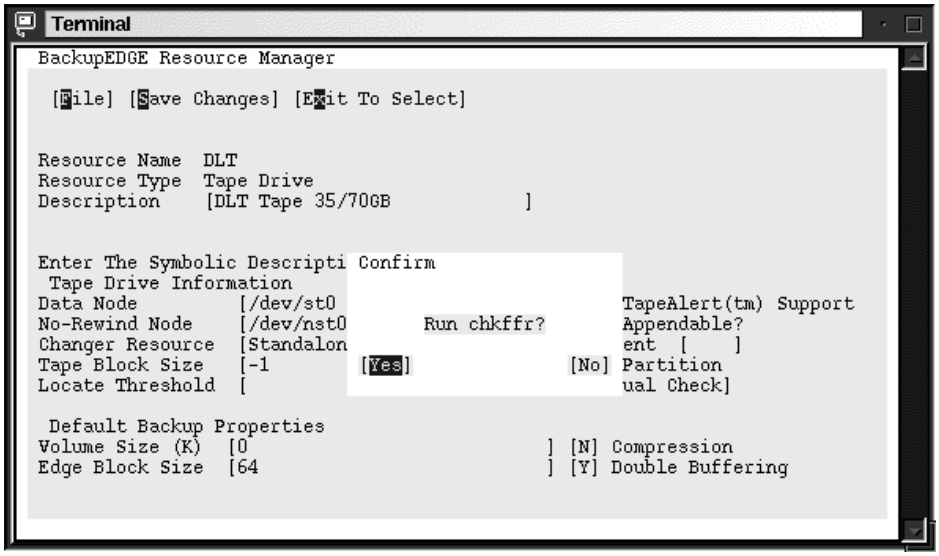

Figure 252. Starting the test

Stop

Performing this test will destroy all data on the tape.

7. Select **Yes** to continue. You will see a window similar to Figure 253.

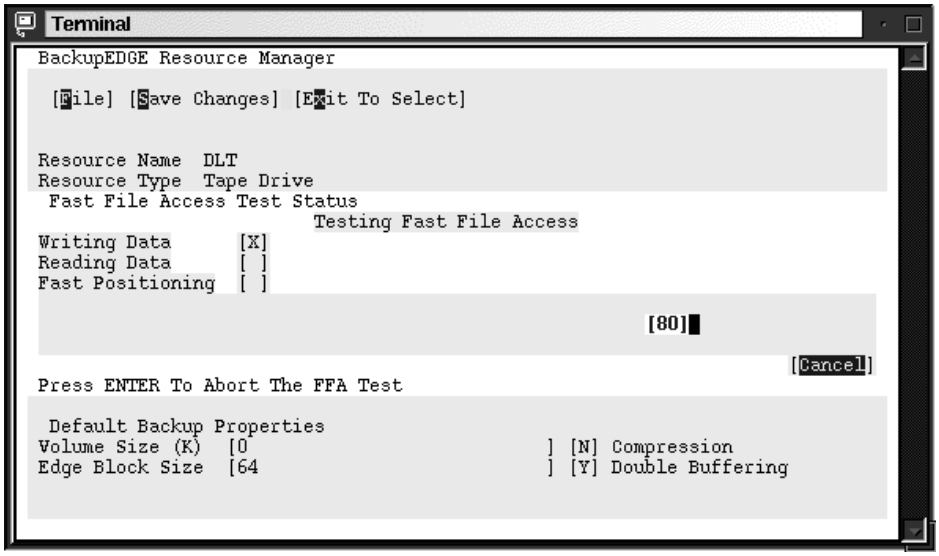

Figure 253. Performance test

After the test is done you will see a window similar to Figure 254.

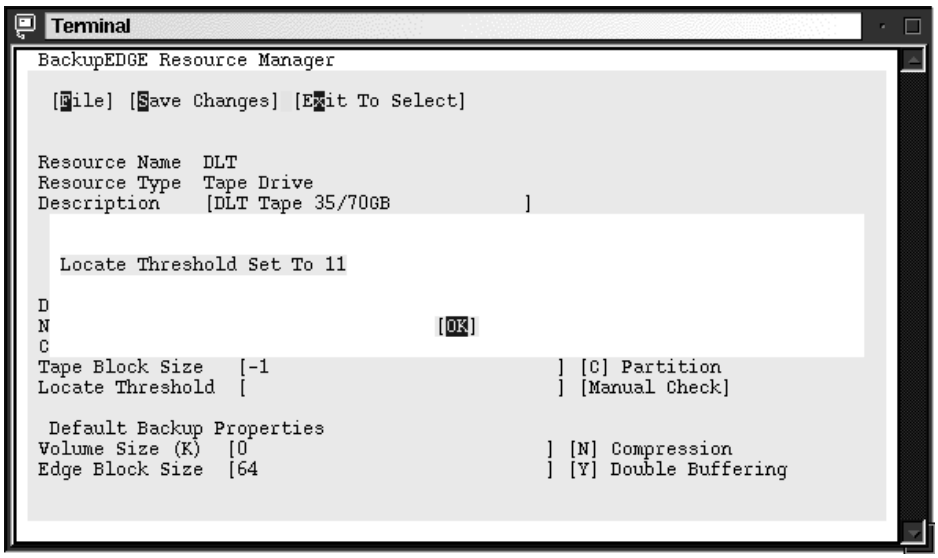

Figure 254. Threshold value

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 251 on page 326. 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 255.

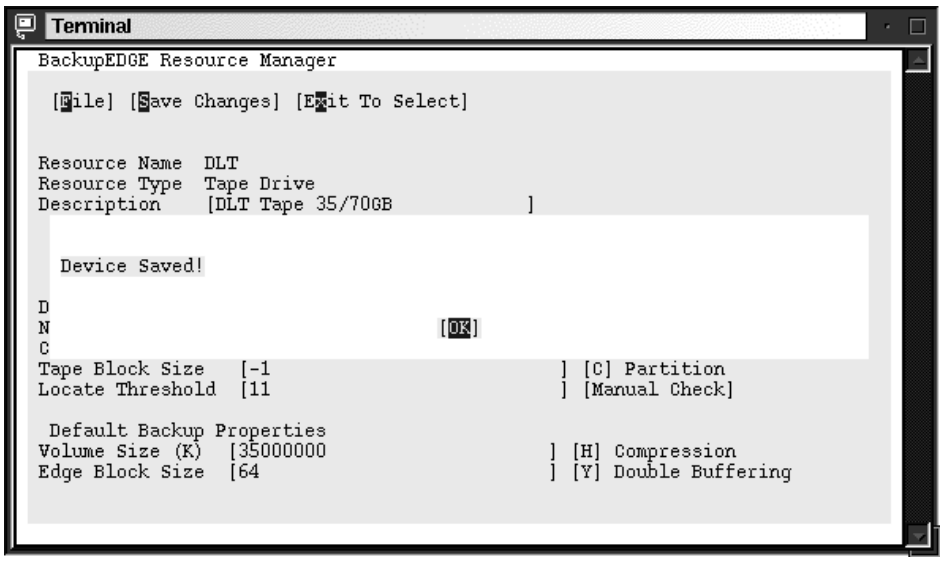

Figure 255. Saving the device definitions

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 256.

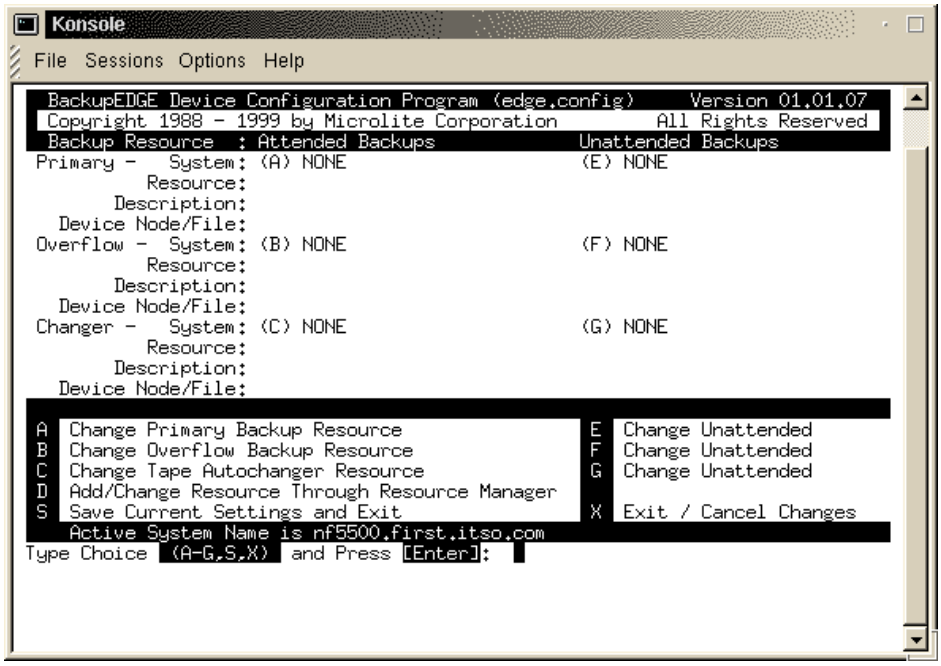

Figure 256. Device Configuration

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 257.

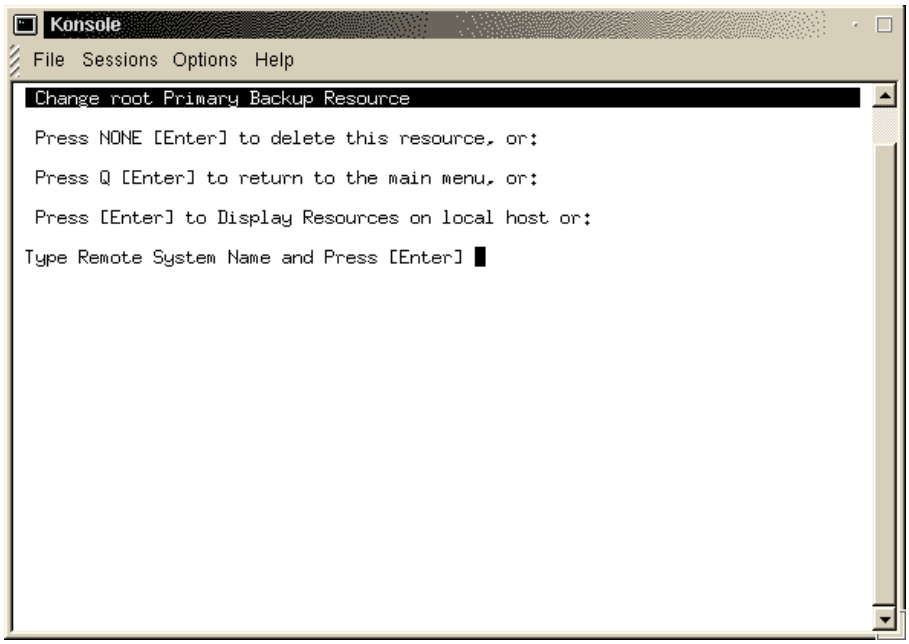

Figure 257. Selecting the device for backup

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 258.

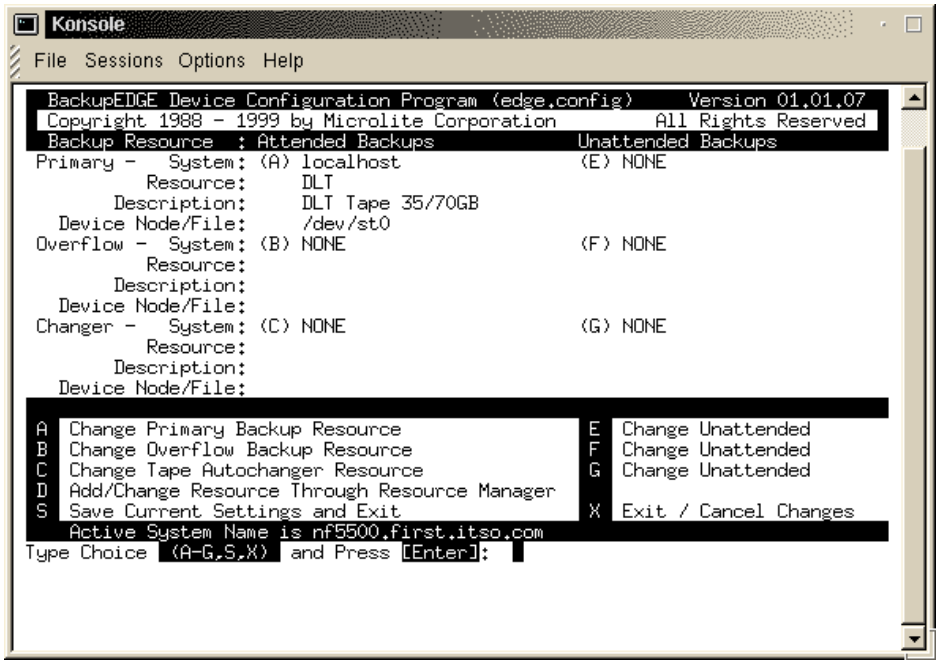

Figure 258. After definition of attended backup device

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 259.

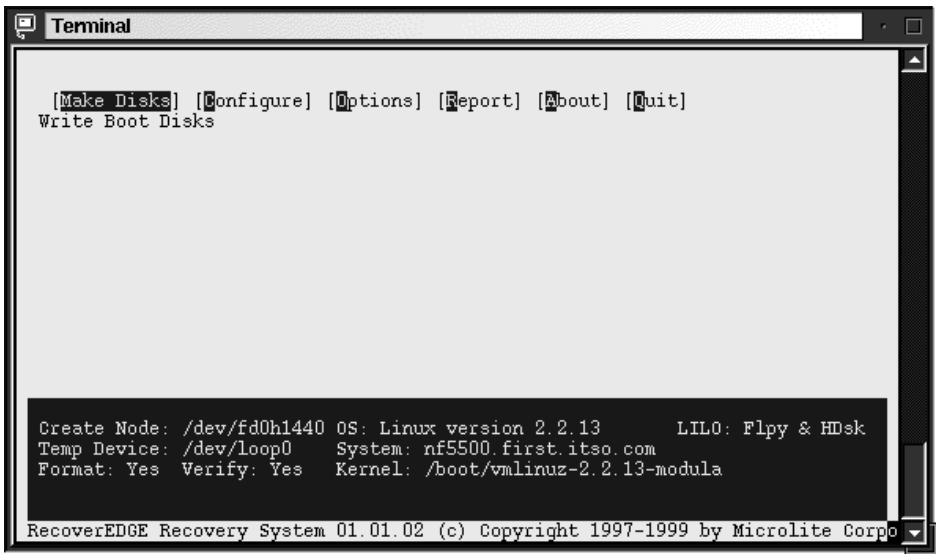

Figure 259. RecoverEDGE utility

2. Select the **Configure** option and press Enter, and you will see a window similar to Figure 260.

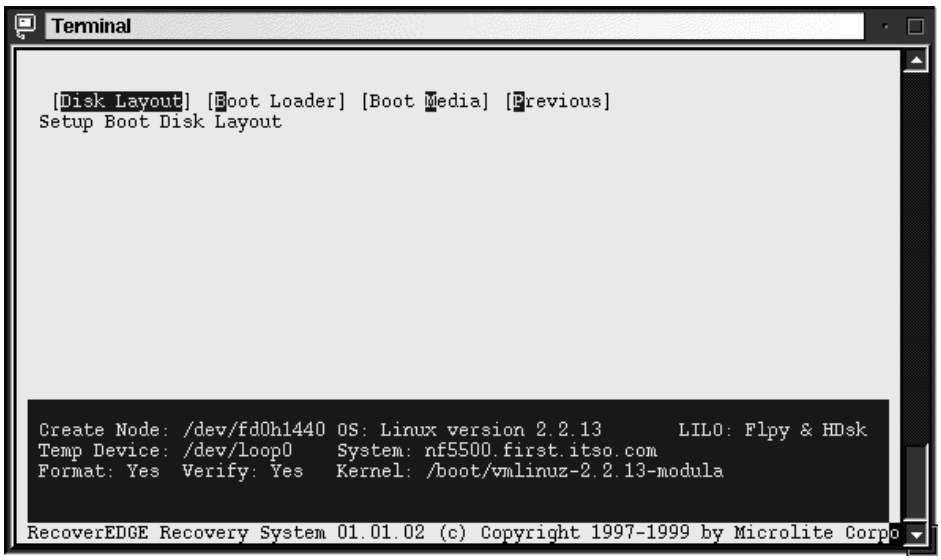

Figure 260. Configure menu

3. Select the **Disk Layout** option and press Enter, and you will see a window similar to Figure 261.

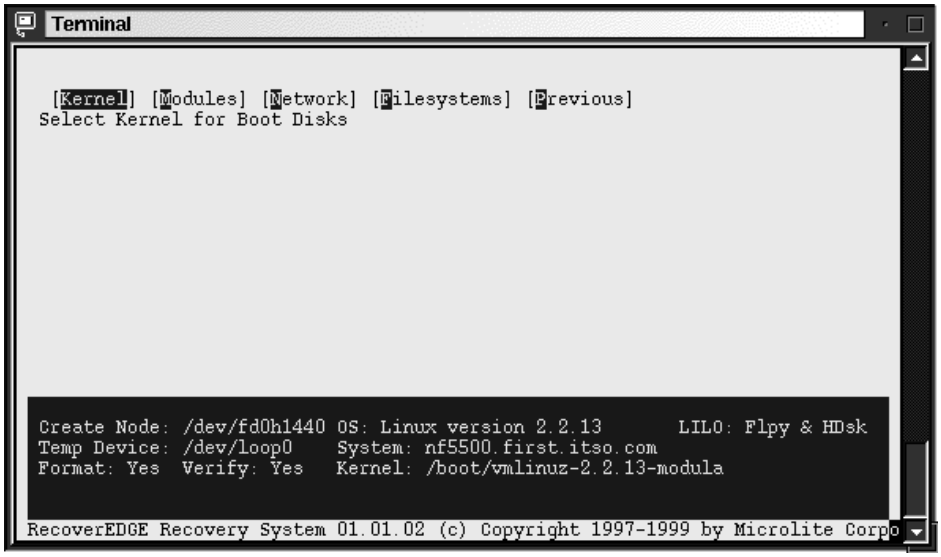

Figure 261. Disk layout menu

4. 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 262.

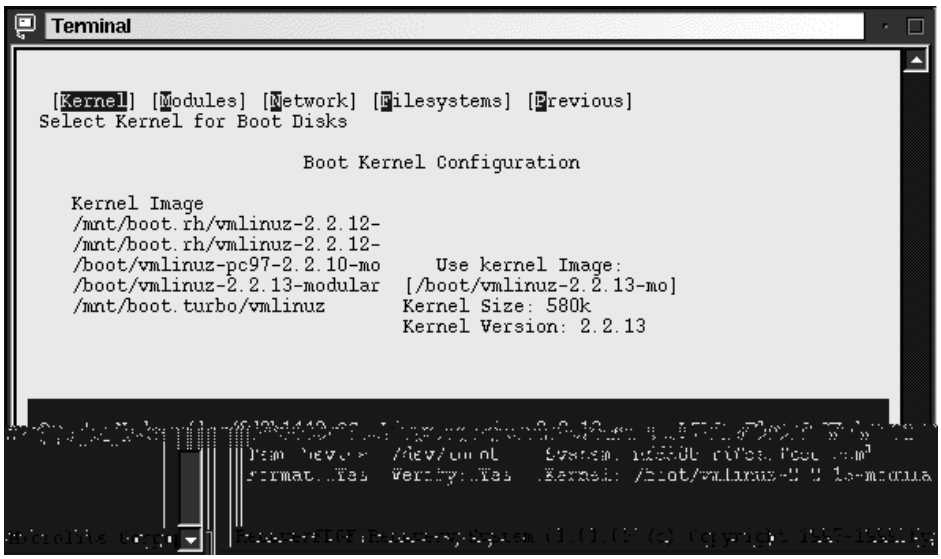

Figure 262. Kernel options

Here you define which kernel will be used for creating the diskette.

5. Return to the previous stage and select **Modules** and press Enter, and you will see a window similar to Figure 263.

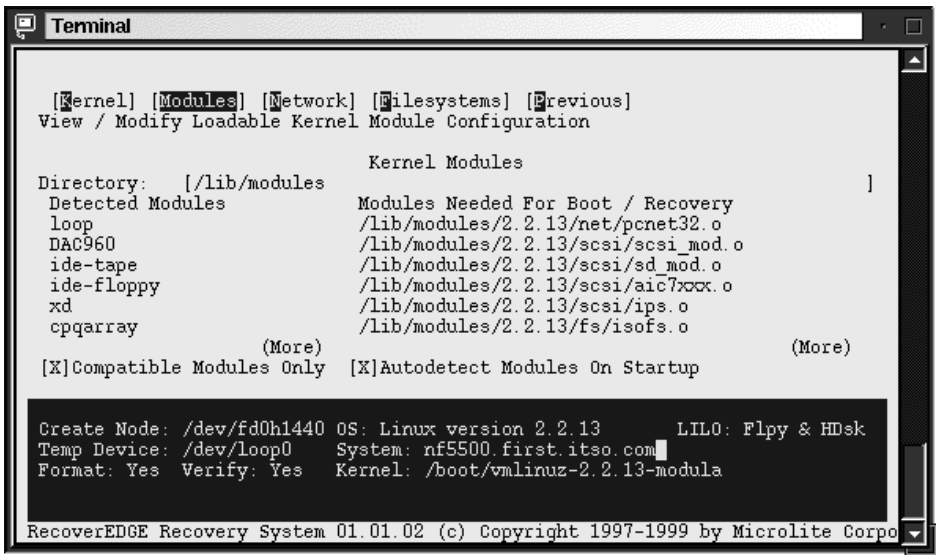

Figure 263. Modules options

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.

6. Return to the previous stage and select **Network** and press Enter, and you will see a window similar to Figure 264.

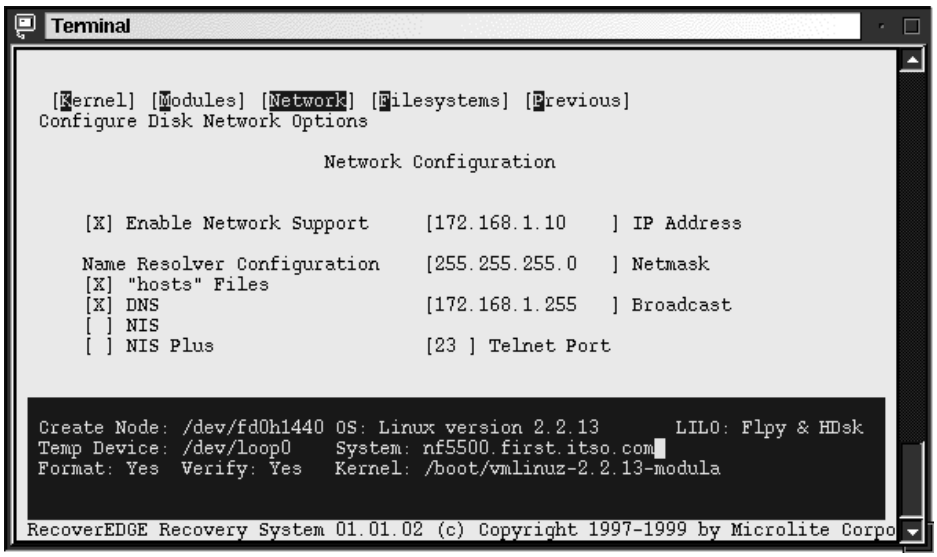

Figure 264. Network options

Here you define you 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 265.

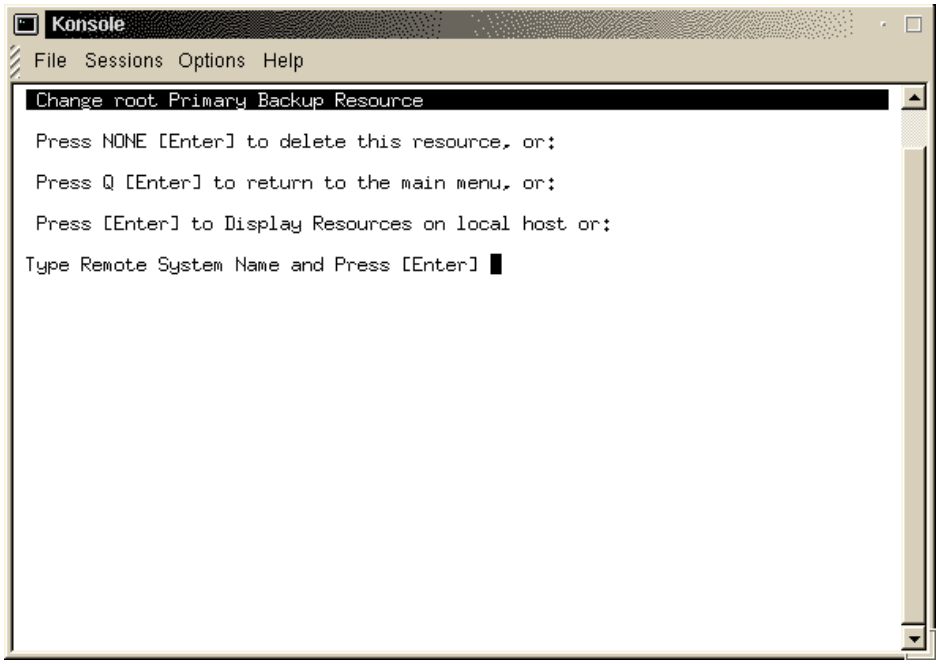

Figure 265. Filesystems options

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 266.

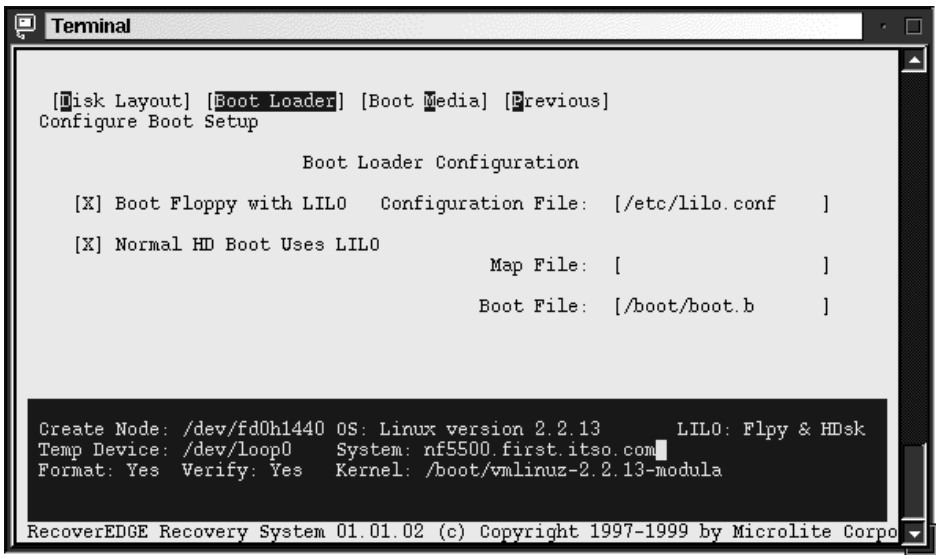

Figure 266. Boot Loader options

Here you define options for the Boot Loader.

9. Return to the configuration panel and select the **Boot Media** option and press Enter. You will see a window similar to Figure 267.

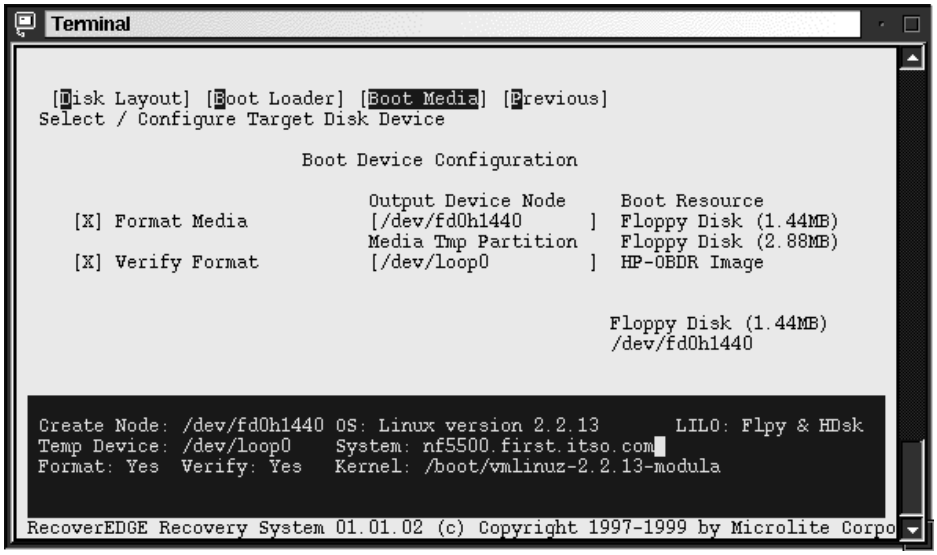

Figure 267. Boot Media options

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 RecoverEGDE boot diskette. After the diskette is loaded, RecoverEDGE will start and you will see a window similar to Figure 268.

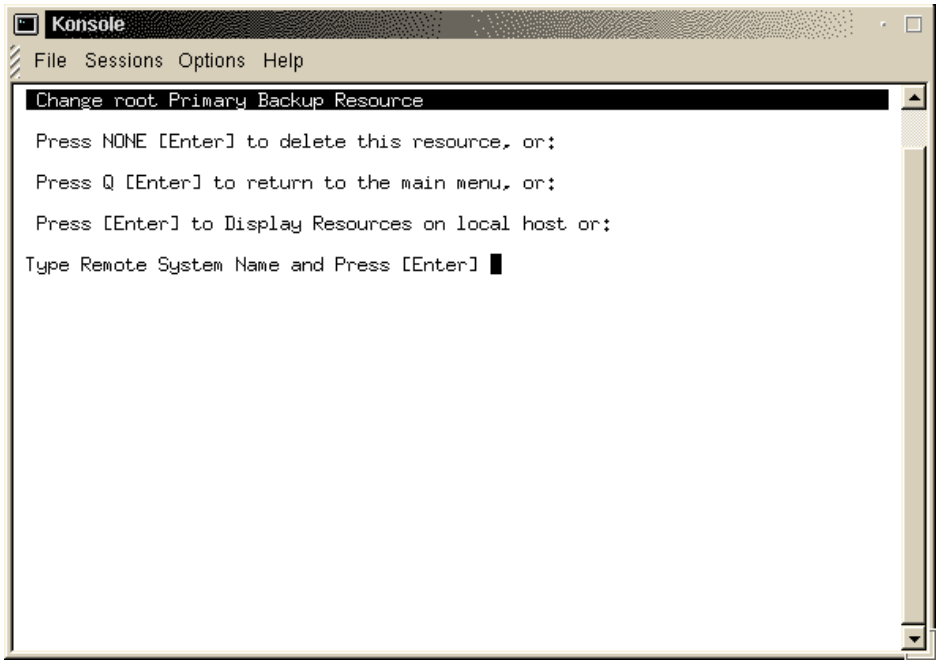

Figure 268. RecoverEDGE initial window

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, your have X-Windows running. Then enter on the command line:

```
Arkeia
```

You will see a dialog like Figure 269:

Figure 269. 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 270):

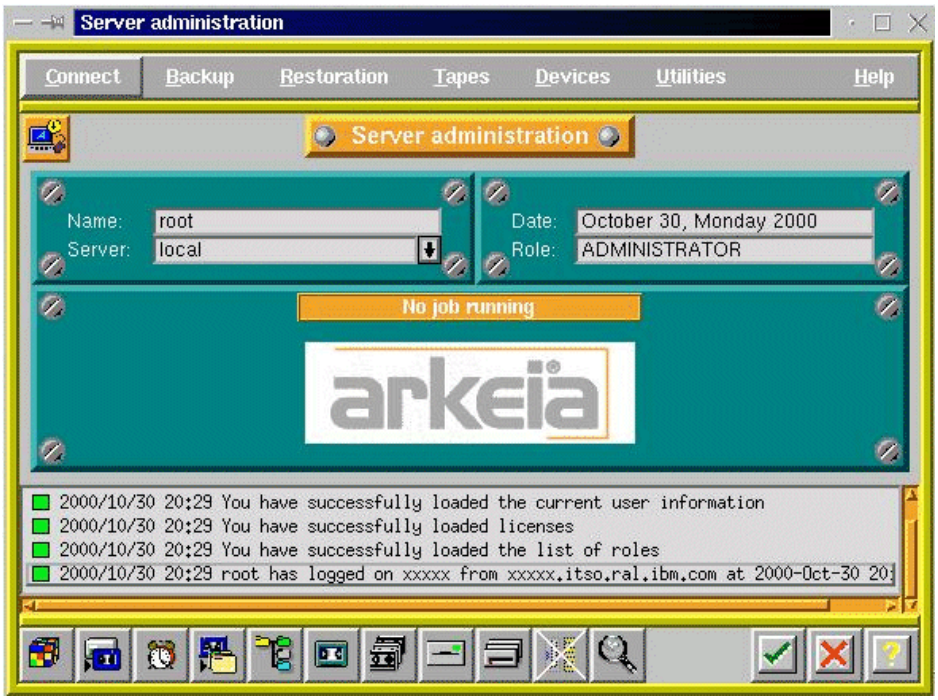

Figure 270. The Arkeia main dialog window

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 271.

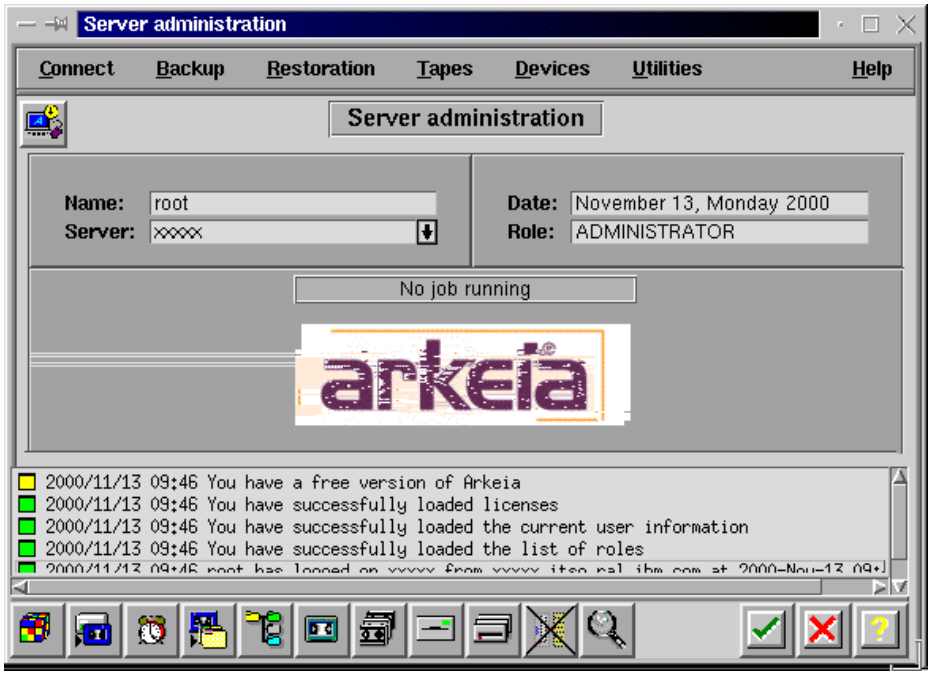

Figure 271. The new Arkeia main dialog window

At the bottom of the window you see push buttons shown in Figure 272:

Figure 272. 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 273:

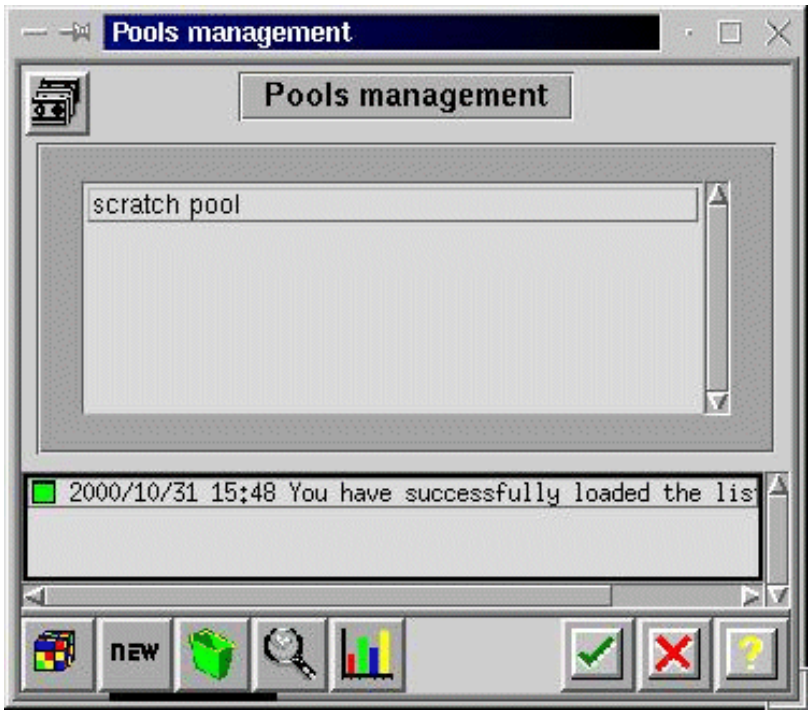

Figure 273. Pools management main dialog window

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 274:

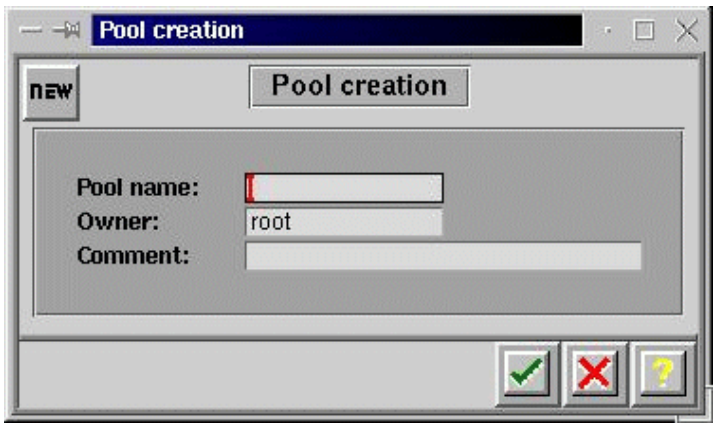

Figure 274. Pool creation window

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 275:

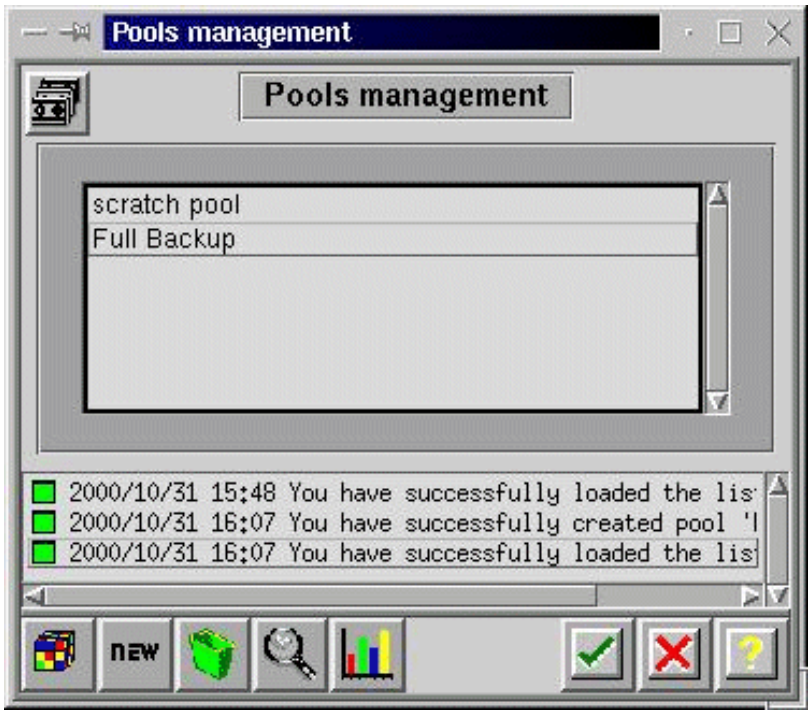

Figure 275. Pools management main window with updated pool list

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 276):

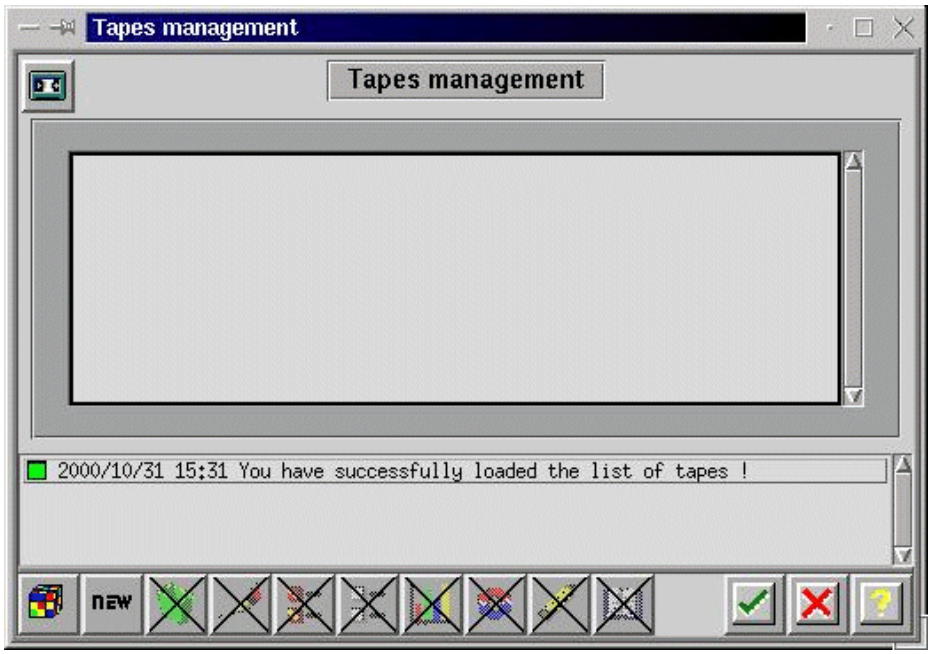

Figure 276. Tape management main window

Click the **new** button to enter new tapes (Figure 277):

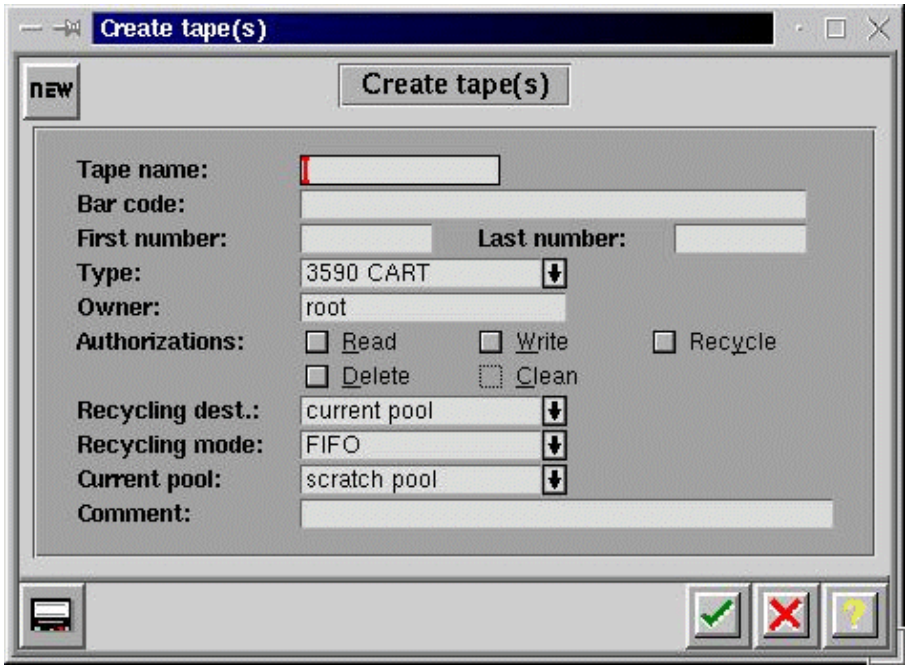

Figure 277. Create tape(s) window

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 278):

Figure 278. 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 279):

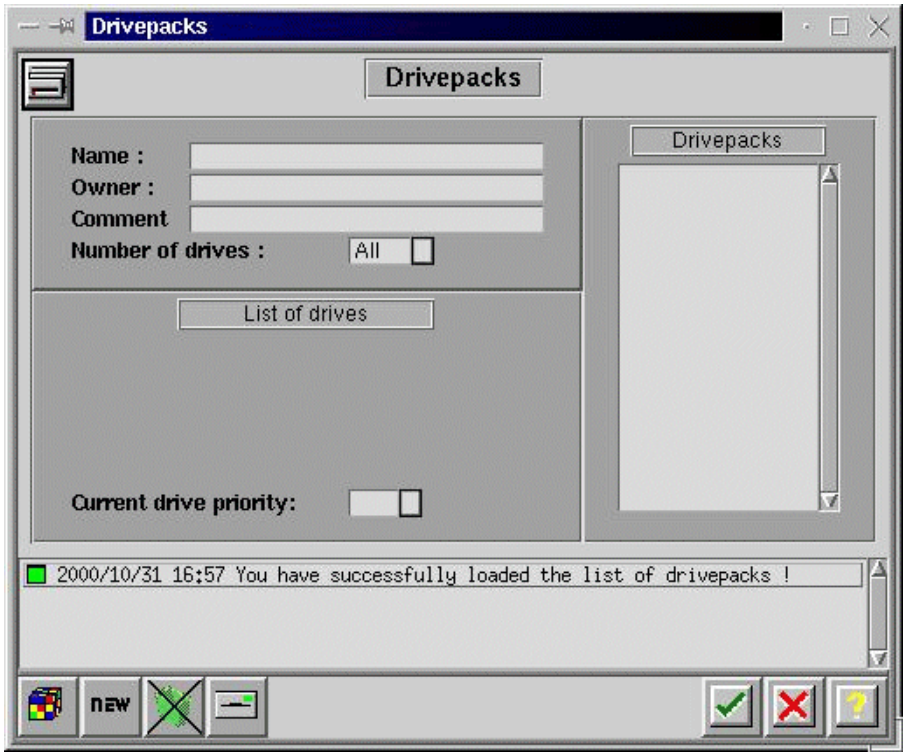

Figure 279. Drivepacks management window

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 280).

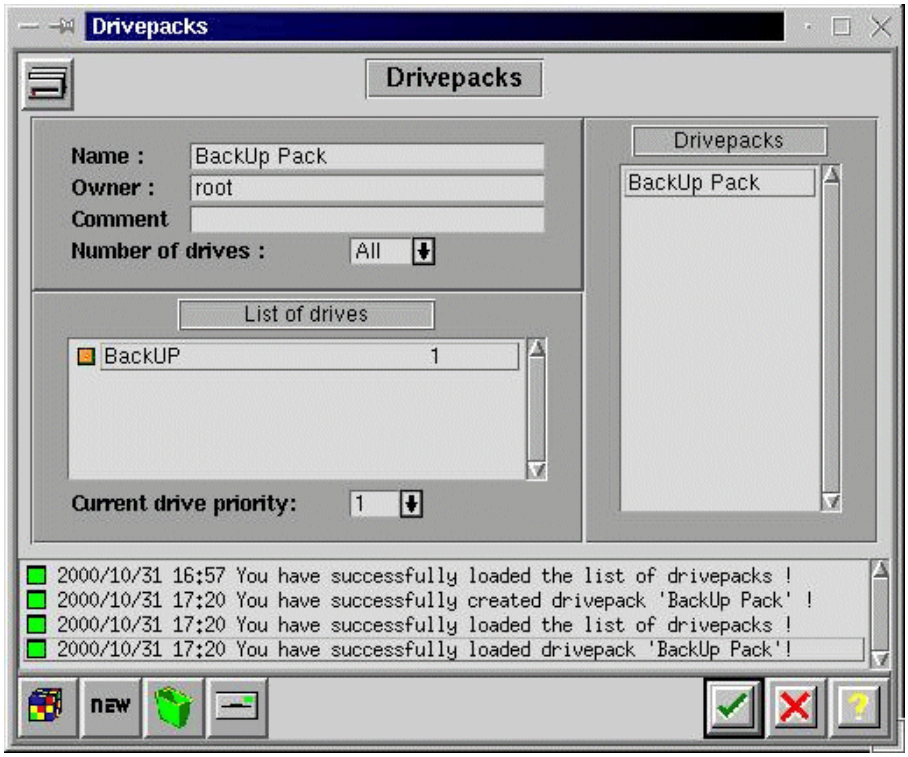

Figure 280. Updated drivepacks management window

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 281:

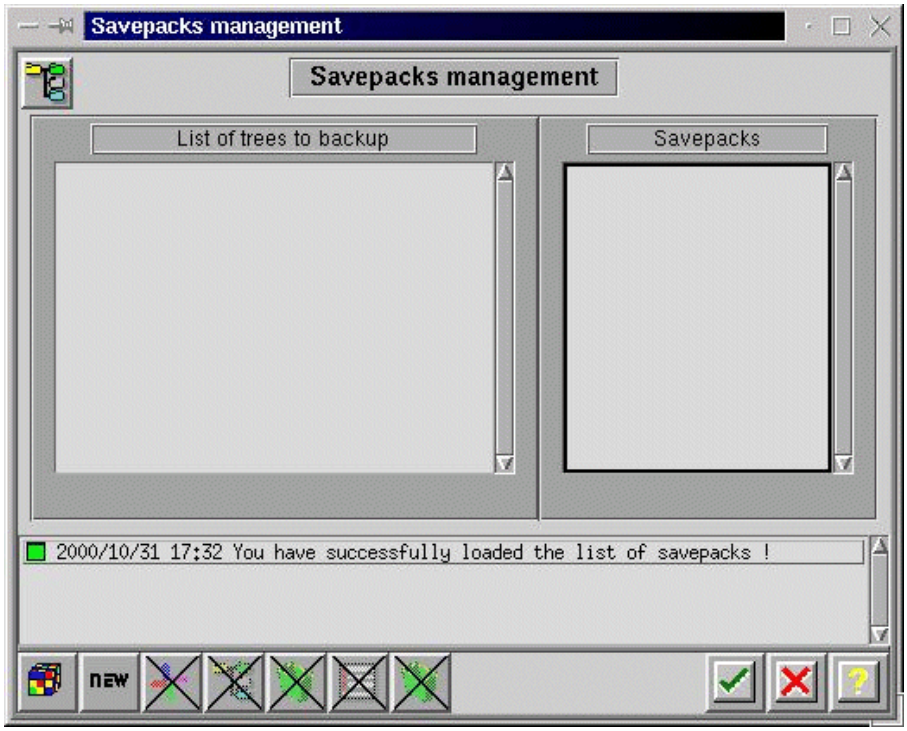

Figure 281. Savepacks management window

Click the **new** button to enter input mode. A window similar to Figure 282 appears.

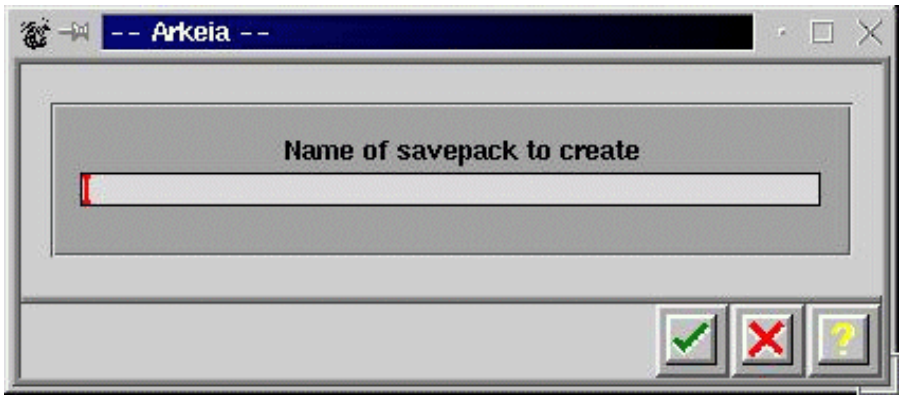

Figure 282. Window to create a new savepack

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 283 appears:

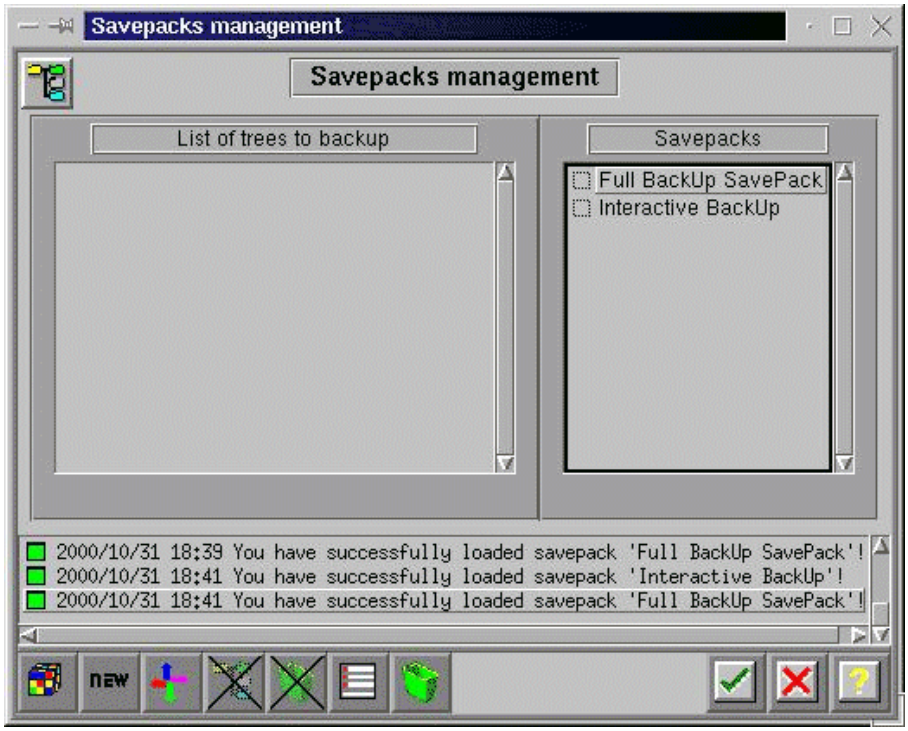

Figure 283. Updated savepacks management window

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 284.

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 285 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.

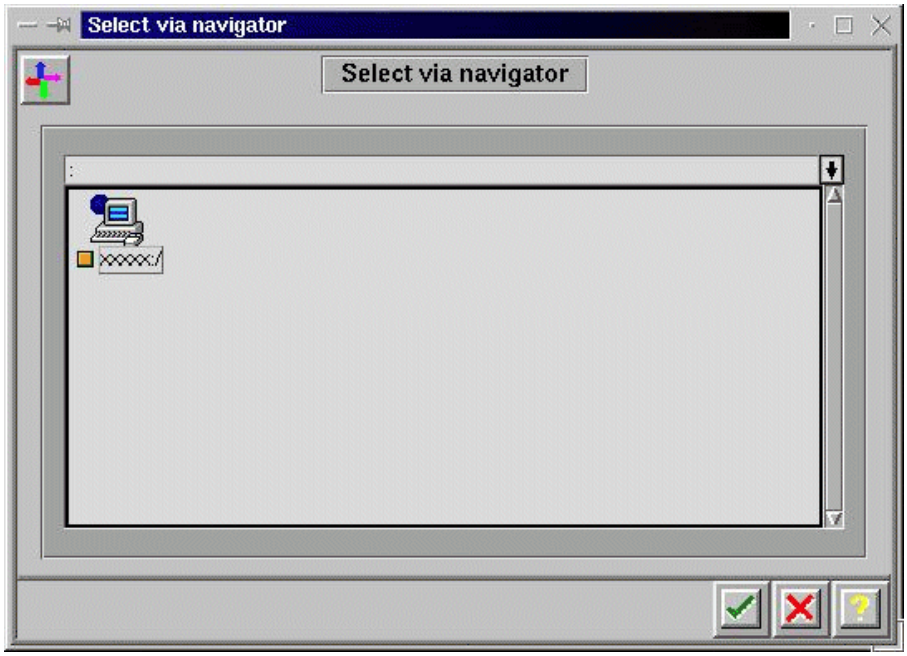

Figure 284. Navigator window

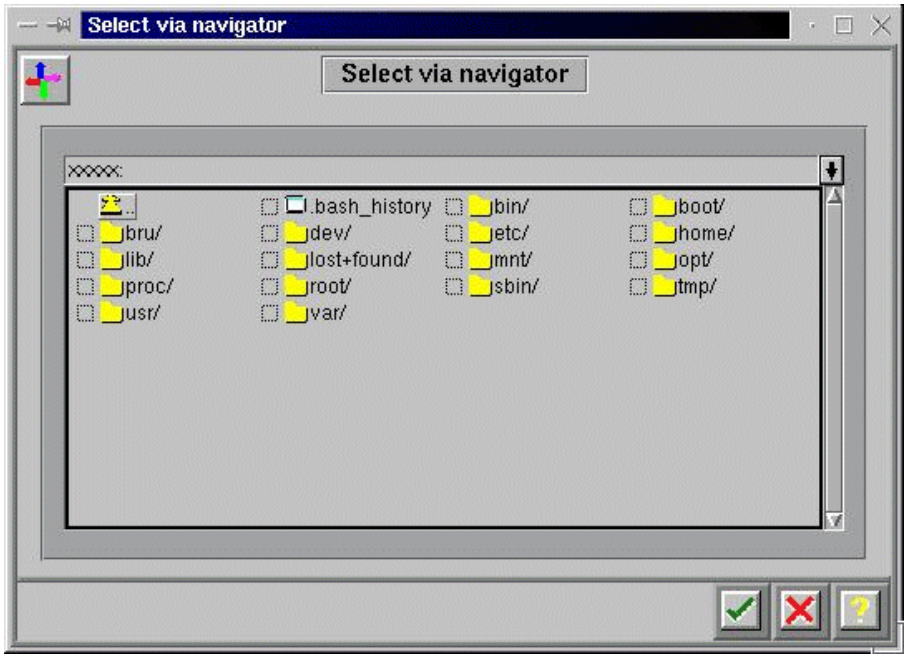

Figure 285. Updated navigator window

To return to the savepacks management window, click the **OK** button. You will see a window similar to Figure 286.

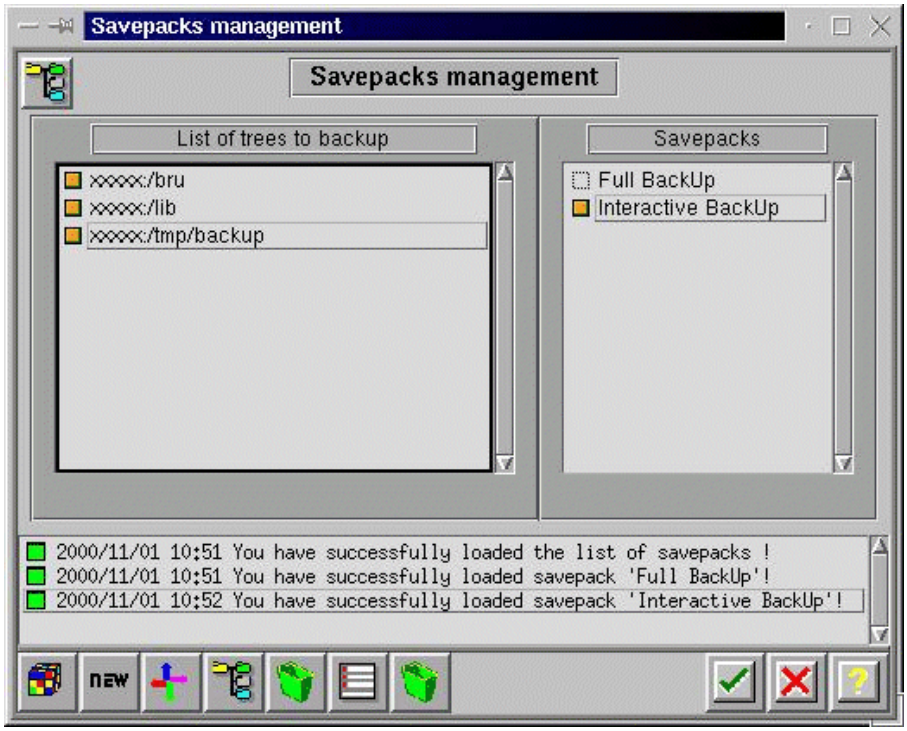

Figure 286. Updated savepacks management window

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 287 appears.

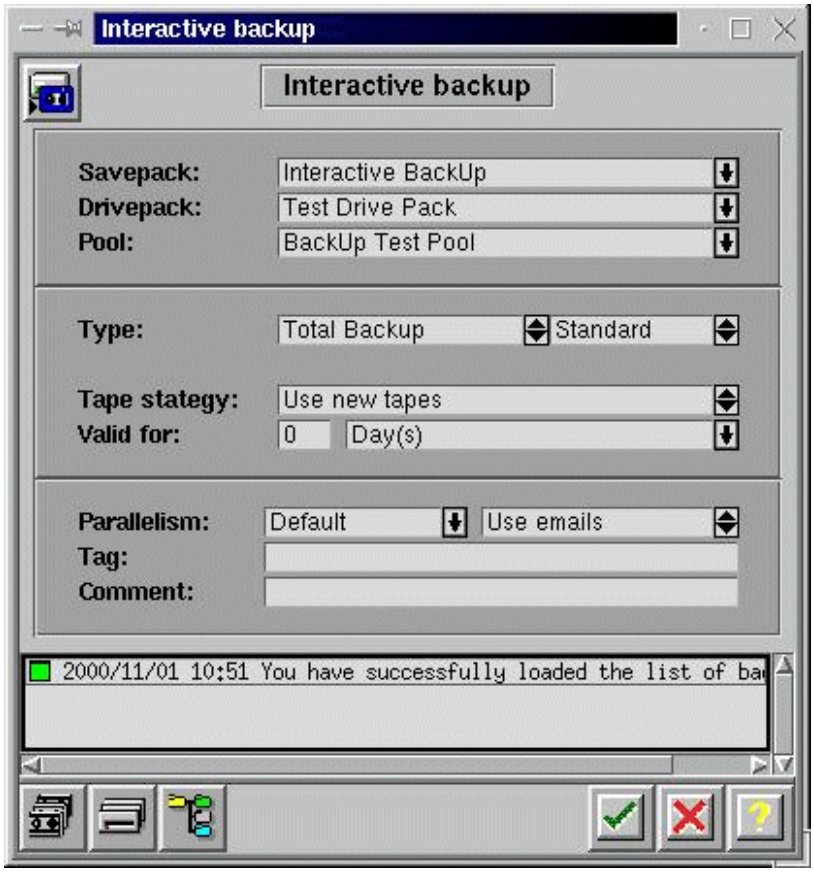

Figure 287. Interactive bckup start window

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 288 appears.

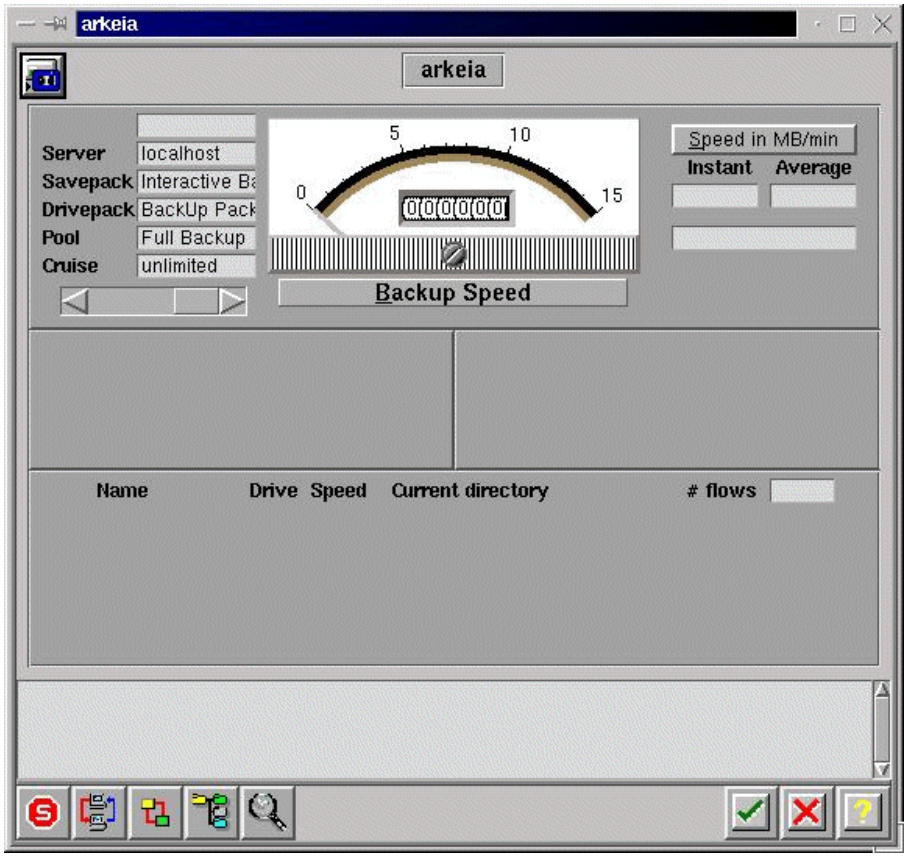

Figure 288. Arkeia's main window during backup

As the backup process proceeds, the content of this window will change. Most of the time, you will see a window like Figure 289.

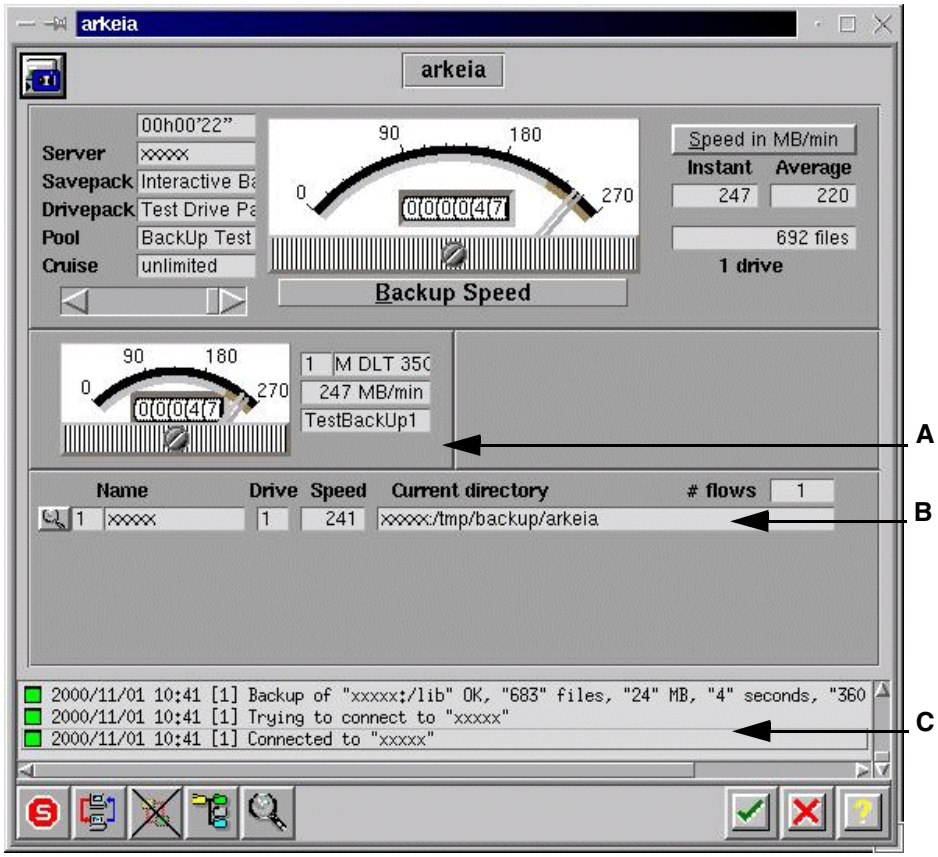

Figure 289. Main window during backup in progress

There are three areas in the window, marked **A**, **B** and **C** in Figure 289, 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 271. 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 290 appears.

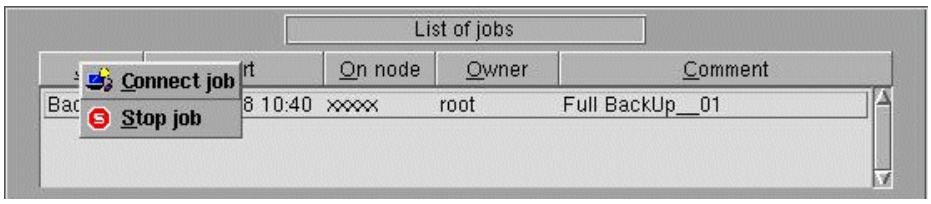

Figure 290. Connect job pull-down menu

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 289 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 291.

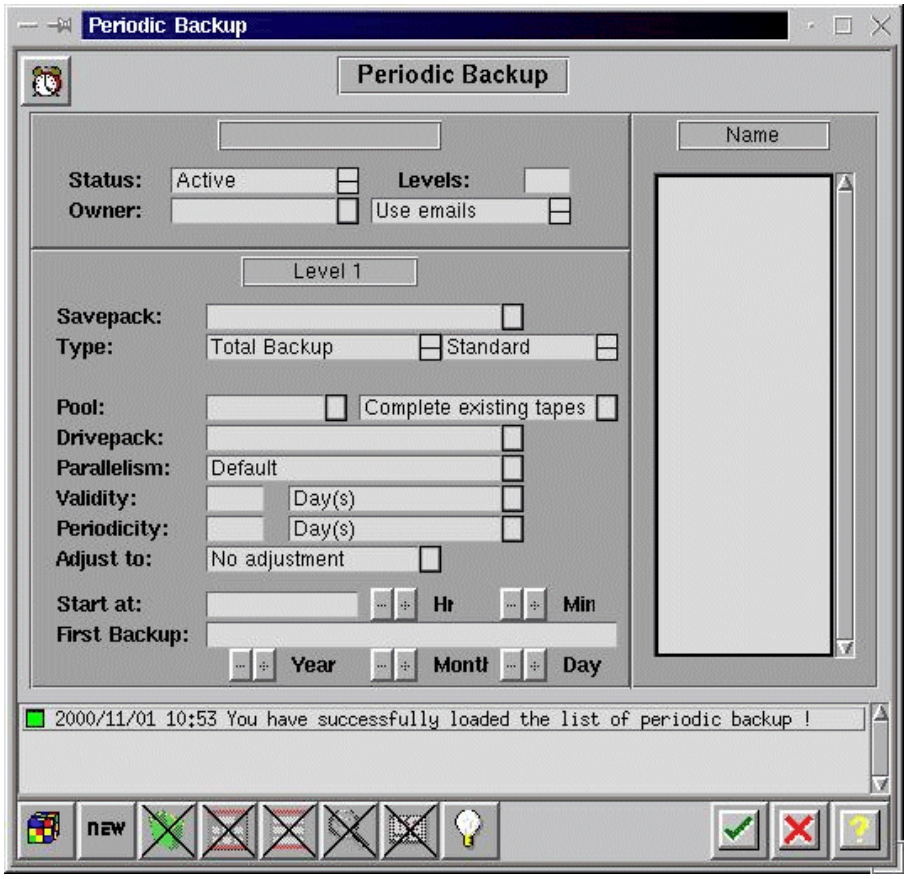

Figure 291. Periodic Backup window

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 292.

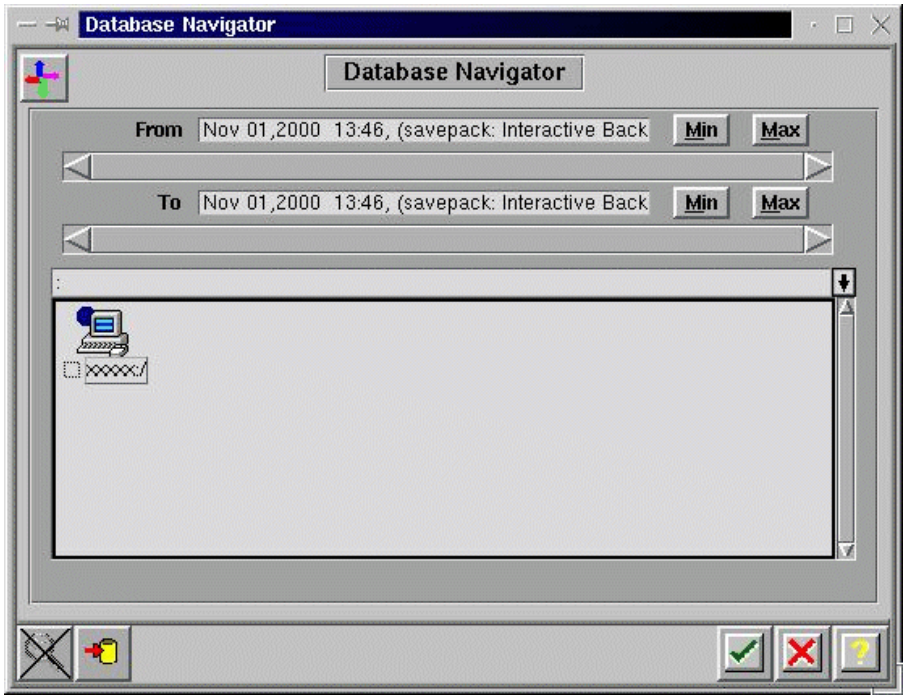

Figure 292. Restoration start dialog

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 293 appears, containing a list of the files or directories that will be restored.

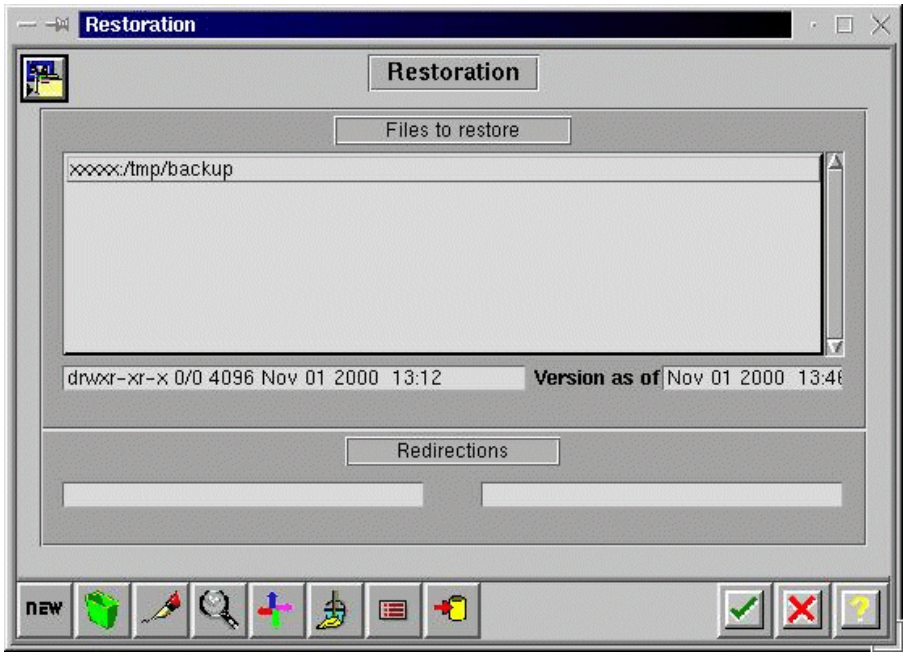

Figure 293. List of directories/files to store

Click the **OK** button in this window opens a new window, shown in Figure 294.

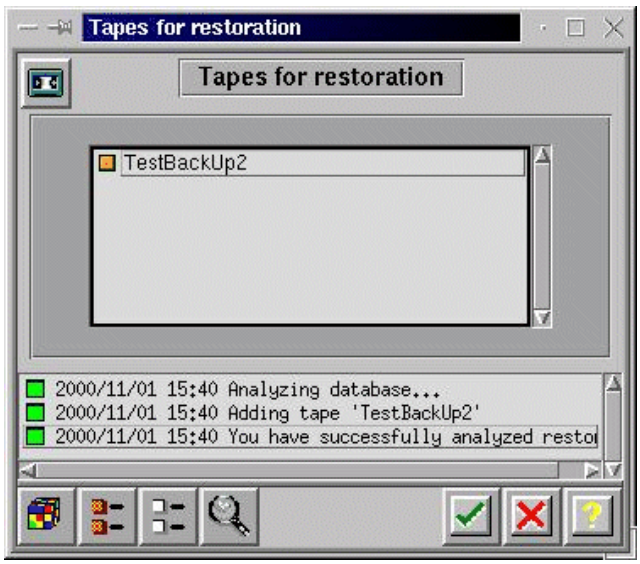

Figure 294. List of tapes used for restoration

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 295.

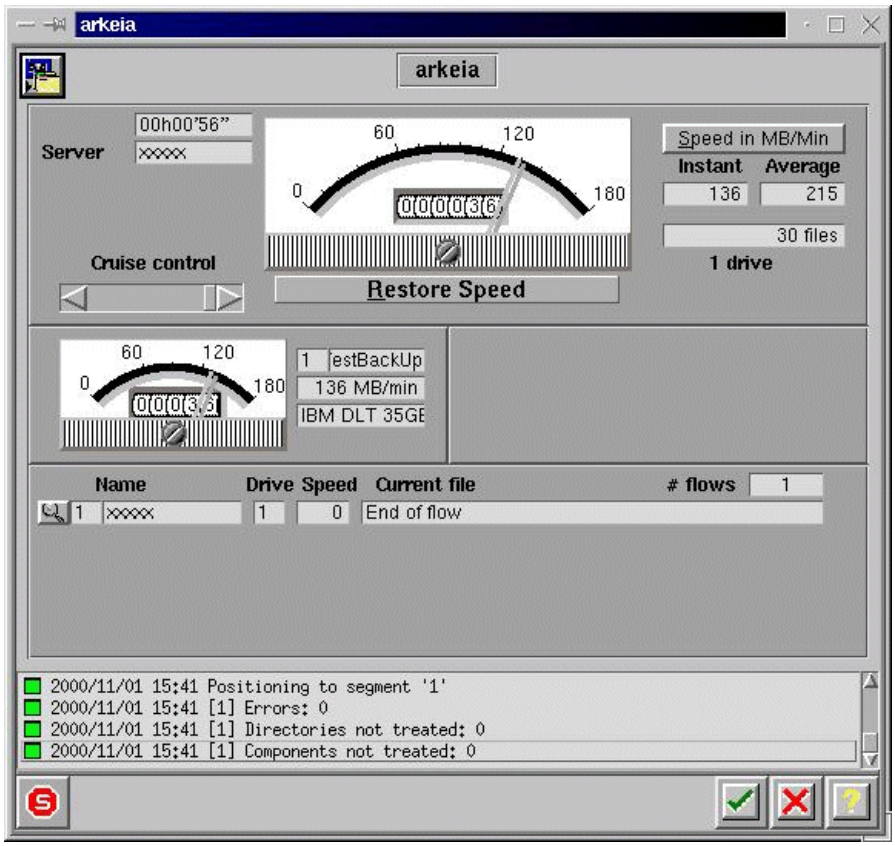

Figure 295. Restoration's main window

If the tape to start with must be mounted, a window like Figure 296 appears.

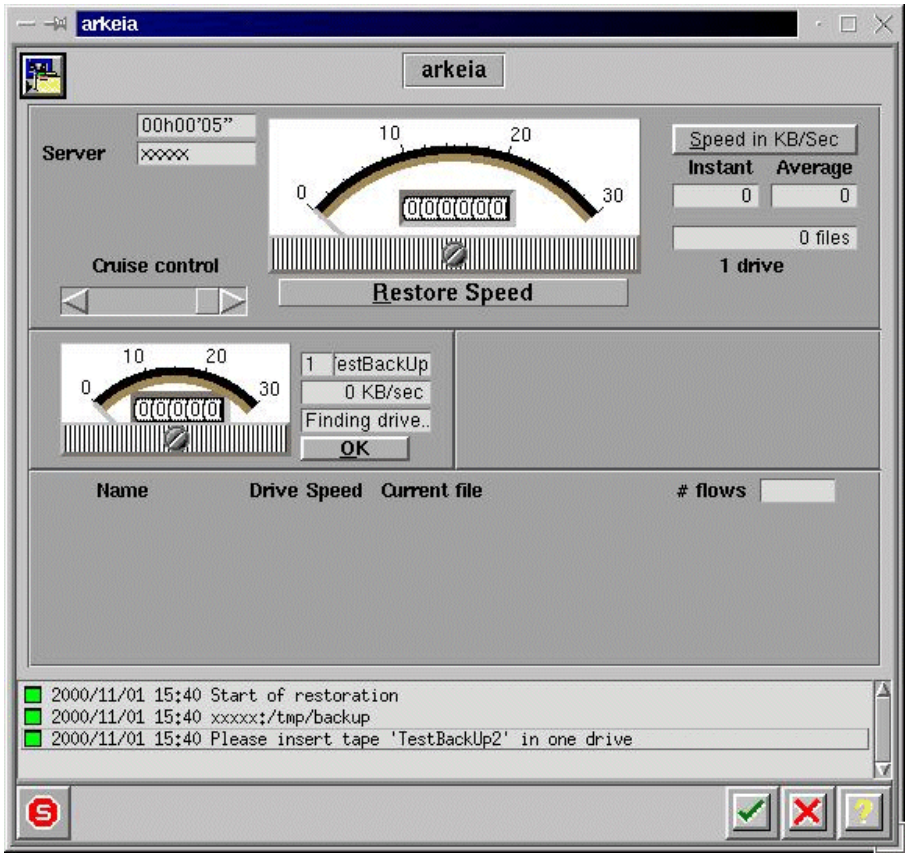

Figure 296. Window during restoration if manual intervention is required

Perform the action required and click **OK** to proceed. The appearance of the window changes. It is now like Figure 295.

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
```

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 377.

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 297:

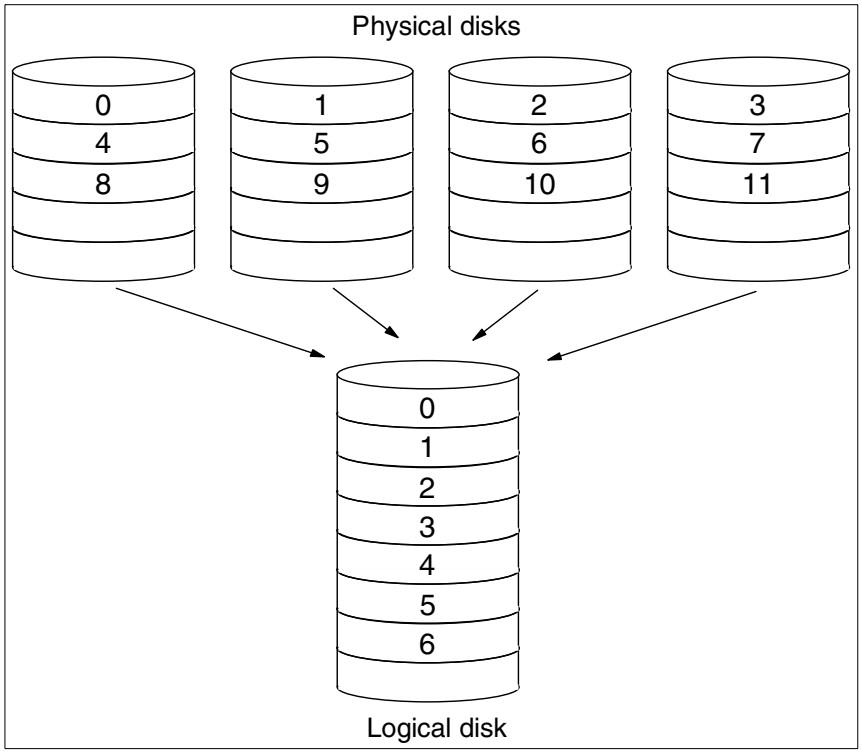

Figure 297. RAID-0 implementation

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 298:

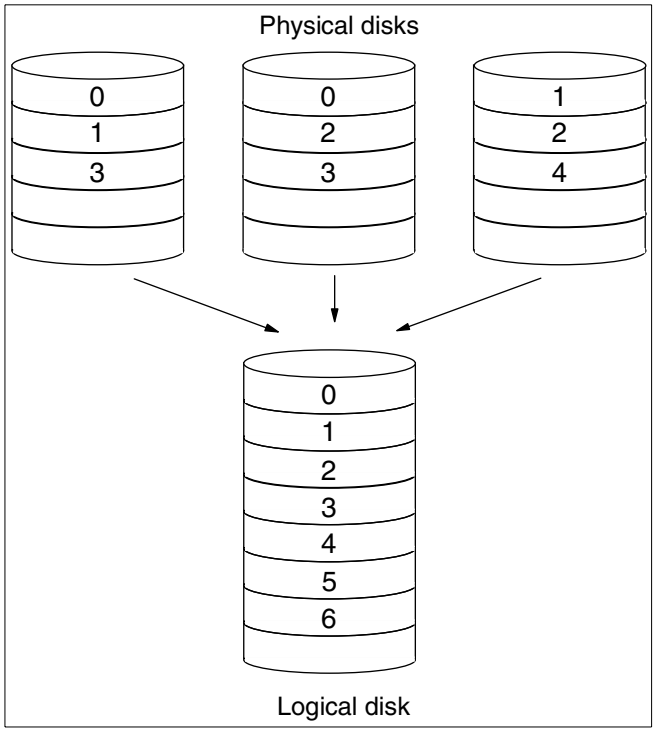

Figure 298. RAID-1E implementation

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:

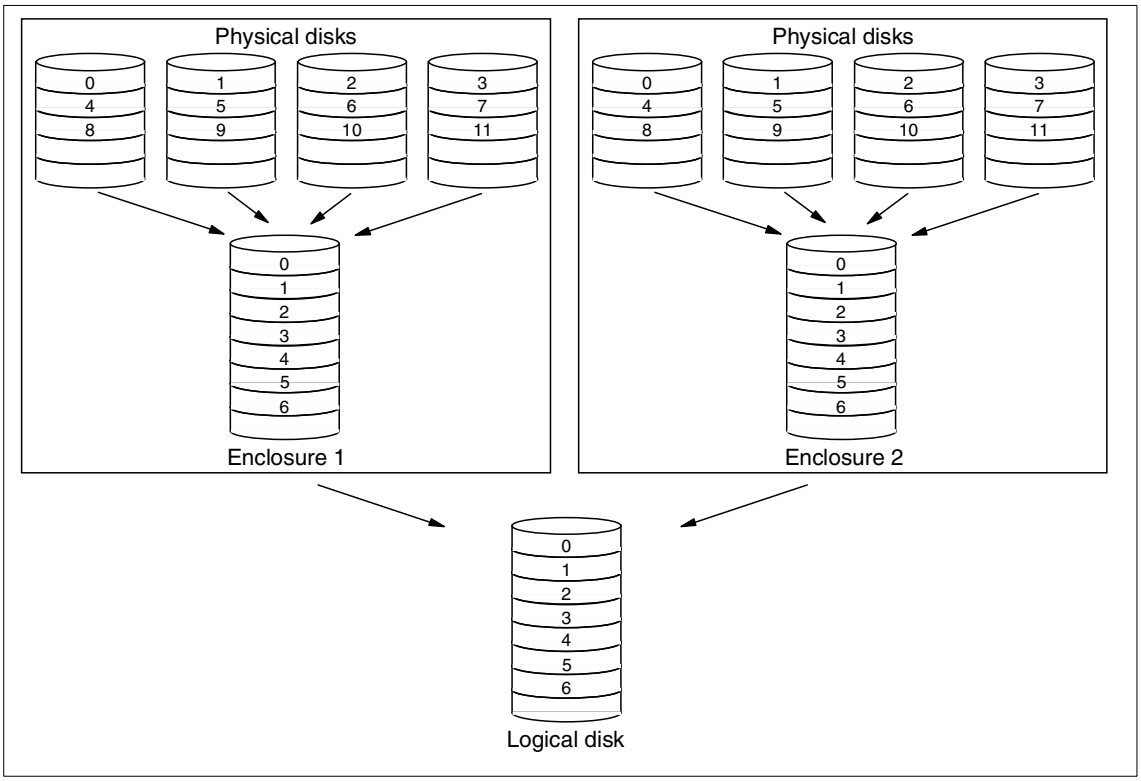

Figure 299. RAID-10 configuration

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 300, spreading the load across the entire array.

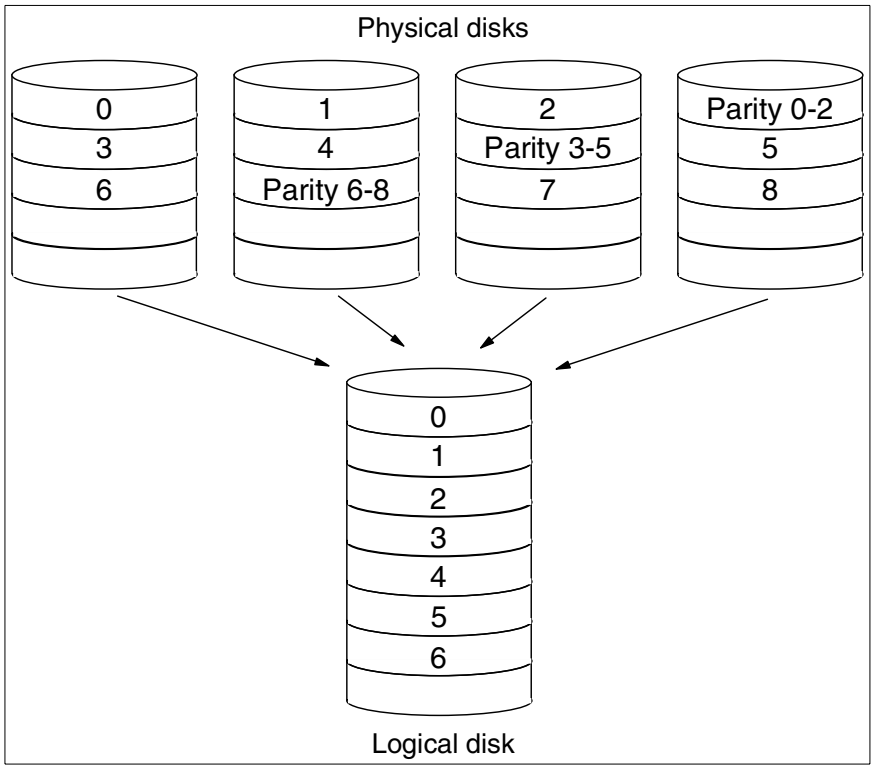

Figure 300. RAID-5 implementation

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 301 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:

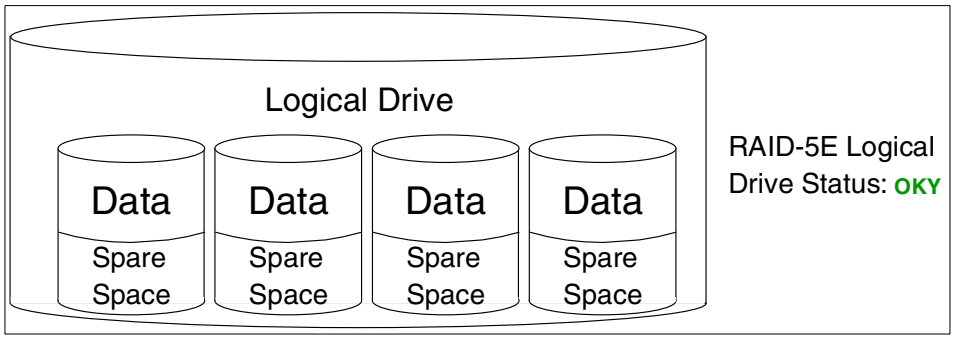

Figure 301. RAID-5E array: normal operation

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.

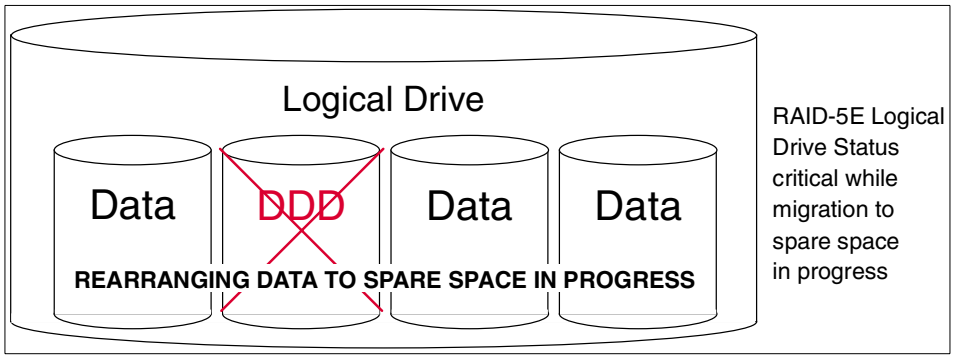

Figure 302. RAID-5E array: single physical disk failure

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 303.

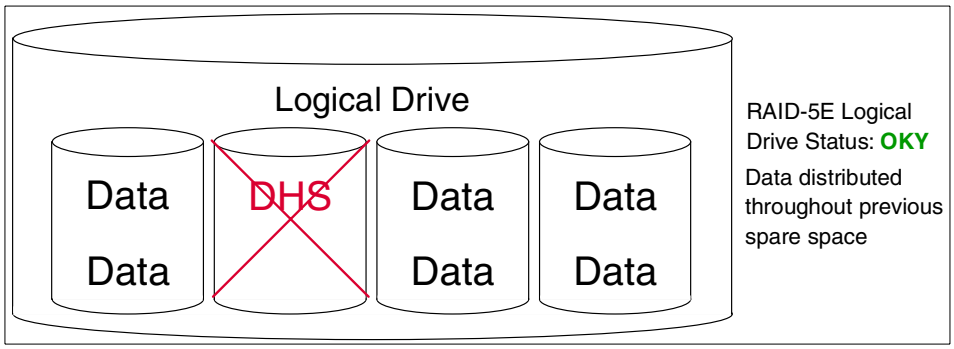

Figure 303. RAID-5E array: data distributed throughout previous spare space

In the event of a second physical disk failure before the previously failed disk has been replaced, illustrated in Figure 304, 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 373.

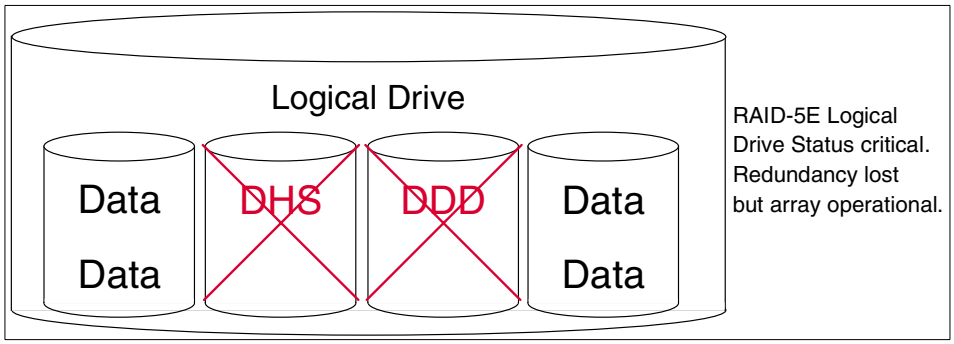

Figure 304. RAID-5E array: second physical disk failure

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 305 on page 381, speed is a significant issue when deciding on RAID level. The numbers shown in this figure and in Figure 306 on page 382 are based on benchmark testing performed by the xSeries and Netfinity server development team. Specific systems may not show precisely the same performance ratios but the figures are representative of typical performance data.

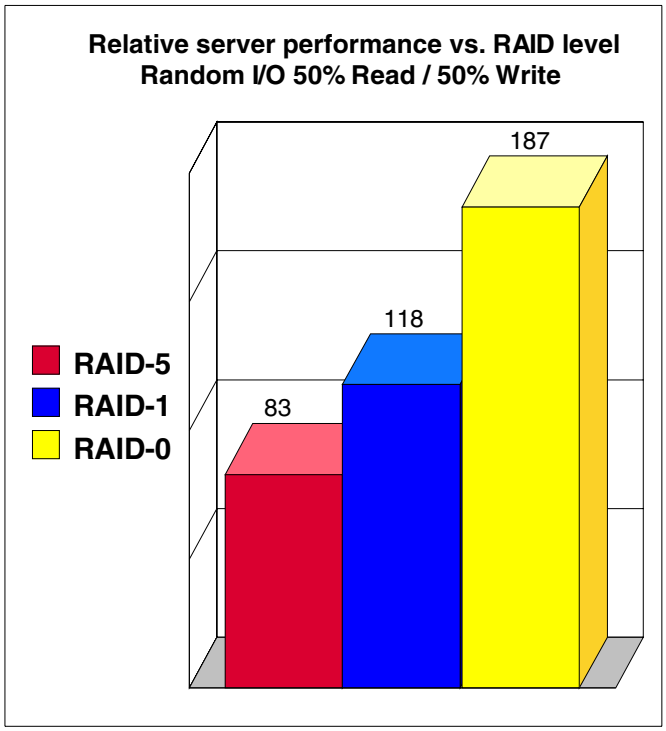

Figure 305. Relative server performance versus RAID strategy

It is important to point out that the speed difference in Figure 305 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 306 on page 382 shows what happens with I/O throughput when we add drives to a RAID-0 array.

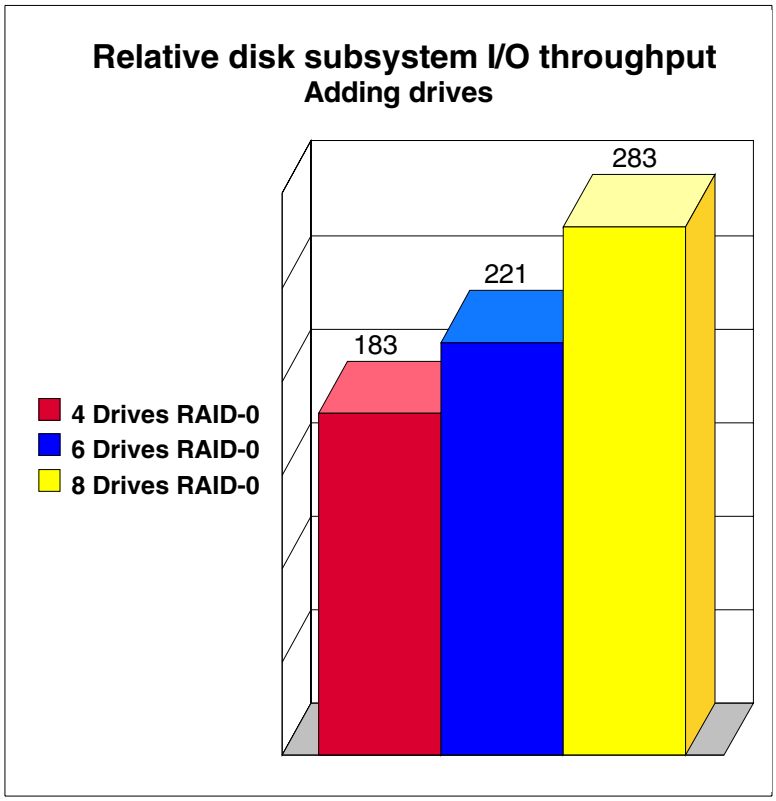

Figure 306. Adding drives to an array

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 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. Working video modes for IBM Netfinity servers

In this appendix you can find some working modes for the Xfree86 servers for IBM Netfinity servers. All working graphics cards in IBM Netfinity servers use the XFree86 SVGA server. So before you do a video card probe, select the SVGA server. These are the modes tested in our working environment. If you want, you can try to find another working mode by yourself. This can be a time-consuming job. Good luck!

- Netfinity 3000

  Modeline "1152x864/70Hz" 92 at 24bpp

- Netfinity 3500M10

  Modeline "1024x768/70Hz" 75 at 16bpp

- Netfinity 5000

  Modeline "800x600/72Hz" 50 at 16bpp

- Netfinity 5600

  Modeline "800x600/85Hz" 60.75 at 24bpp

- Netfinity 5500 M10

  Modeline "1024x768/70Hz" 75 at 8bpp

- Netfinity 5500 M20

  Use XFree86 VGA server

- Netfinity 7000 M10

  Use XFree86 VGA server

- Netfinity 8500R

  Modeline "1152x864/70Hz" 92 at 24bpp

Appendix C. Recommendations for disk partitions

In Table 25 you can see the recommended values for the partition sizes.

Table 25. Suggested partition and file system scheme

File system	Minimum (in MB)	Recommended (in MB)
/boot	10	25
swap	=RAM	=RAM multiply by two
/	50	100
/usr	1000	1500
/var	25	128
/tmp	1	128
/opt	1	512
/home	5	2048

Appendix D. Hardware issues for IBM Netfinity servers

In this appendix we explain what you need to consider when you install Caldera OpenLinux on IBM Netfinity servers. This is because Caldera OpenLinux 2.3 and Caldera OpenLinux eServer have some limitations when installing on IBM Netfinity servers. Below you can find some hints for installing Caldera OpenLinux on IBM Netfinity servers:

1. Netfinity 3000
   - Although S3 Trio3D graphics chip is supported in XFree86 3.3.4, which is included in Caldera OpenLinux 2.3, and XFree 3.3.5, which is included in eServer, only some modes are working.
2. Netfinity 3500M10
   - S3 Savage4 graphics chip is not supported in XFree86 3.3.4 included in Caldera OpenLinux 2.3, so you can only use VGA server.
   - S3 Savage4 graphics chip is supported in XFree86 3.3.5 included in Caldera OpenLinux eServer, but only some modes are working.
3. Netfinity 5000
   - Caldera OpenLinux 2.3 install does not recognize the CD-ROM, so it is impossible to install from CD-ROM.
   - S3 Trio64V2 GX graphics chip is supported in XFree86 included in Caldera OpenLinux eServer, but only some modes are working.
4. Netfinity 5600
   - S3 Trio64 3D graphics chip is supported in XFree86 3.3.4 included in Caldera OpenLinux 2.3 and in XFree86 3.3.5 included in Caldera OpenLinux eServer, but only some modes are working.
   - AMD Am79C973 onboard Ethernet adapter is not supported in kernels up to 2.2.12.
5. Netfinity 5500 M10
   - S3 Trio64V2 GX graphics chip is supported in XFree 3.3.4 included in Caldera OpenLinux 2.3 and in XFree 3.3.5 included in Caldera OpenLinux eServer, but only some modes are working.
6. Netfinity 5500 M20
   - S3 Trio64V2 GX graphics chip is supported in XFree 3.3.4 included in Caldera OpenLinux 2.3 and in XFree 3.3.5 included in Caldera OpenLinux eServer, but because of the implementation on planar board only VGA server is working.

7. Netfinity 7000 M10
   - S3 Trio64V2 GX graphics chip is supported in XFree 3.3.4 included in Caldera OpenLinux 2.3 and in XFree 3.3.5 included in Caldera OpenLinux eServer, but only some modes are working. Because of the implementation on planar board, only the VGA server is working.
8. Netfinity 8500R
   - S3 Trio3D graphics chip is supported in XFree 3.3.4 included in Caldera OpenLinux 2.3 and in XFree 3.3.5 included in Caldera OpenLinux eServer, but only some modes are working.

Appendix E. 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 = LINUXRULZ

# server string is the equivalent of the NT Description field
   server string = Samba Server on Red Hat Linux

# 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, lprng, aix, hpux, qnx
   printing = lprng

# 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 oneWINS 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 F. Modified ifup-dhcp file

```
#!/bin/sh

function saveFile {

   file=$1
   save=$1.save
   dhcp=$1.dhcp

   test -f $dhcp || return 0
   test -f $save || cp $file $save
   mv $dhcp $file
   return 0
}

function restoreFile {

   file=$1
   save=$1.save

   test -f $save && mv $save $file
   return 0
}

interface=$1

case $0 in
*ifup*)
   echo "Trying to obtain network configuration via DHCP."

#  /sbin/dhcpcd -HD $interface >/dev/null || {
   /sbin/dhclient $interface >/dev/null || {
      echo "No response from DHCP server" >&2
      exit 1
   }

#  if test -f /var/run/dhcpcd-$interface.info; then
#     . /var/run/dhcpcd-$interface.info
#     if test -n "$NISDOMAIN"; then
#        cat >/etc/nis.conf.dhcp <<-EOF
#        #
#        # Generated by DHCP boot script
#        #
#        EOF
#        if test -n "$NISSERVER"; then
```

```
#           echo "domainname $NISDOMAIN server $NISSERVER"
#        else
#           echo "domainname $NISDOMAIN broadcast"
#        fi >> /etc/nis.conf.dhcp
#     fi
#  fi

   saveFile /etc/nis.conf
   saveFile /etc/resolv.conf
   : ;;
*ifdown*)
   echo "Shutting down DHCP daemon."
   /sbin/dhcpcd -k $interface

   restoreFile /etc/nis.conf
   restoreFile /etc/resolv.conf
   : ;;
*)
   echo "$0: unknown mode of operation for DHCP script" >&2
   exit 1;;
esac

exit 0
```

Appendix G. Using the additional material

This redbook also contains additional material in diskette format, and Web material. See the appropriate section below for instructions on using or downloading each type of material.

G.1 Using the diskettes

The diskettes that accompanies this redbook contains the following:

File name	*Description*
boot.msg	MSG (message) file
initrd.gz	GZ file
ldlinux.sys	System file
syslinux.cfg	CFG (configuration) file
vmlinuz.sys	File

G.1.1 System requirements for using the diskettes

The following system configuration is recommended for optimal use of the diskettes:

Hard disk space:	none (MB minimum)
Operating System:	Caldera OpenLinux 2.3
Processor:	Intel Pentium III or higher
Memory:	64 MB
Other:	CD-ROM drive

G.1.2 How to use the Web material

You can also access the contents of the diskettes by pointing your Web browser at the Web site ibm.com/redbooks and accessing Additional Materials under the subjects Redbooks. The contents of the diskettes are located in the directory named the same as the redbook form number (SG24-5861). Alternatively, you can create a subdirectory (folder) on your workstation and copy the contents of the diskettes into this folder.

Diskette 1: "Caldera 2.3 - Install diskette with ServeRAID 4.40.03 (supports ServeRAID 2/3/4)". Use this diskette with the installation instructions in this redbook.

Diskette 2: "Caldera 2.3 - Install diskette for Updated CD with ServeRAID 4.00.03". This diskette is used with the Caldera installation CD that includes updated code for some Netfinity servers that already have the ServeRAID

driver 4.00.03 included. Therefore, the installer does not need to use the installation procedure stated in this redbook.

Appendix H. Special notices

This publication is intended to help you install and configure Caldera OpenLinux on IBM @server xSeries and Netfinity servers. The information in this publication is not intended as the specification of any programming interfaces that are provided by Caldera OpenLinux or IBM @server xSeries and Netfinity. See the PUBLICATIONS section of the IBM Programming Announcement for more 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
@server	Redbooks Logo 
Netfinity	Netfinity
NetVista	Lotus
ServeRAID	Lotus Notes
xSeries	Domino

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 Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Linux is a trademark of Linus Torvalds.

Caldera, the C-logo, OpenLinux, and DR-DOS are either registered trademarks or trademarks of Caldera Systems, Inc.

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 I. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

I.1 IBM Redbooks

For information on ordering these publications see “How to get IBM Redbooks” on page 411.

- *Linux for WebSphere and DB2 Servers,* SG24-5850
- *Red Hat Linux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5853
- *SuSE Linux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5863
- *TurboLinux Integration Guide for IBM @server xSeries and Netfinity*, SG24-5862

I.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

I.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.

I.4 Referenced Web sites

These Web sites are also relevant as further information sources:

- http://www.caldera.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/webservers/httpservers/
- http://www-4.ibm.com/software/webservers/httpservers/doc/v136/readme_httpserver.htm
- http://www.apache.org/docs/misc/perf-tuning.html
- http://www-4.ibm.com/software/webservers/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.ibmlink.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.ibmlink.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.ibmlink.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.ibmlink.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.

IBM Redbooks fax order form

Please send me the following:

Title	Order Number	Quantity

First name	Last name	
Company		
Address		
City	Postal code	Country
Telephone number	Telefax number	VAT number

☐ Invoice to customer number _____

☐ Credit card number _____

Credit card expiration date	Card issued to	Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

A

accounts 99
  managing accounts 101, 106
Adaptec SCSI Controller 4
Apache 199
  features 199
  installation 201
  performance tips 209
Arcnet drivers 127
array 58

B

backup 297
backup and restore utility 297
BackupEDGE
  unattended operation 307
basic system administration 87
  console 88
  KDE 89
  kpackage 93
  login 88
  package
    install 95
    uninstall 94
  RPM 97
  terminal 90
bind4 189
bind8 189
BIOS 4
block devices 112
block interleave data striping 383
BM HTTP Server
  Administration Server 205
booting from CD-ROM 5
BRU 301
  basic backup 299
  basic restore 299
  commands 299
  installation 297
  restore 305
  scheduling 304

C

Caldera Open Administration System (COAS) 29, 87
Caldera OpenLinux 3, 4, 6, 35, 37, 38
  basic system administration 87
  Caldera Open Administration System (COAS) 91
  hardware requirements 3
  installation 3, 5
  KDE windows manager 89
  kpackage 93
  multiple PCI buses install 35
  package management 97
  partitioning 15
  root account 26
  ServeRAID installation 36
  token-ring install 83
  using ServeRAID manager 75
  video setup 10
  Webmin tool 129, 131, 132, 229, 268
  Webmin tools 132
  XF86Setup 129
  X-Windows 90
Caldera Systems 1, 2, 75
CD-ROM 3, 6
CD-ROM drivers 127
character devices 112
COAS 90, 91, 167, 214, 249, 251, 257
  accounts 99
    managing accounts 101
    managing groups 106
  daemons 109
  filesystem 110
  hostname 111
  kernel modules 127
  network menu 120
  peripherials
    mouse 115
    printer 116
  peripherials menu 114
  resources 111
  services 109
  system menu 98
  time 113
  token-ring install 83
Common Internet File System (CIFS) 155
custom setup 15

D

DAP 240

DARPA 258
data Compression 306
data stripe mirroring 383
Defunct Disk Drive (DDD) 377
Defunct Hot Spare (DHS) 70
DHCP 29, 122, 197, 198
  token-ring 85
disk mirroring 371, 383
disk striping 370
disk subsystem
  See also RAID
  RAID performance 380
display adapter 3
distributed hot spare 384
distribution and updates 297
DMA 112
DNS 125, 214
Domain Name System (DNS) 187, 188, 189, 197
  configuration 190
  fully qualified domain name (FQDN) 187
  installation 189
  IP addresses 187
  name resolution 188
drive status 58

E

error detection 297
error recovery 297
Ethernet 29, 121, 123
Ethernet drivers 127
even parity 374
extended partition 18

F

fast response cache accelerator 201
Fibre Channel 369
file comparisons 297
file overwrite protection 297
file/print 155
filesystem permissions 138
firmware 4
format partitions 22
full backups 297

G

group identification
  GID 102

H

hard drives 3
hardware requirements 3
hdparm 153
hostname 111
Hot Spare (HSP) 73
hot spare drives 377
hot spare rebuild 68
hot swap rebuild 61, 62

I

IBM commitment to Linux 1
IBM Development and Competency Centers for Linux 1
IBM HTTP Server 199, 200
  features 200
  installation 202
  performance tips 210
IBM Technology Center 1
IETF 240
incremental backups 297, 306
installation type 25
installing Linux 3, 5
  custom setup 15
  DHCP 29
  hardware issues 389
  installation type 22, 25
  keyboard setup 9
  LILO 30
  logon 35
  monitor setup 12
  mouse setup 8
  network setup 29
  partitions 15, 16
  root password 26
  ServeRAID 36
  Tetris 32
  time zone 31
  token-ring 83
  video setup 10
Internet 2
interrupts 112
intranet 2
introduction 1
IO ports 112
ISDN drivers 128
ISO 240

J
Java runtime 76

K
KDE 89, 146
kernel modules 127
keyboard 9
kpackage 93, 156, 214, 257
  check dependencies 96
  install 95
  replace file 96
  replace package 96
  test 96
  uninstall 94
  upgrade 96
KTop 146

L
LANAID diskette 85
LDAP 200, 239, 241
LDIF 241
LILO 30, 45
Linus Torvalds 1
Linux kernel 1
logical drive 19
logical drive access 58
logical drives 58
logical partition 20
login 88
  console 88

M
manage printers 117
master backups 306
master boot record (MBR) 30
Microlite
  BackupEDGE 306, 308
    backup 310
    features 306, 342
    incremental backup 316
    installation 307
    master backup 316
    master restore 319
    restore 314
    schedule backup 320
    tape device 323
  RecoverEDGE 332
    boot disks 333
    features 332
    total crash 340
mirrored RAID-0 arrays 383
monitor 3, 12
monitor setup 12
mount point 17
mouse 115
mouse setup 8
mouse type 3
multimedia drivers 128
multiple PCI buses 35
multitasking 25
multivolume archives 297

N
name resolution 124
NetBIOS name server (NBNS) 155
Netfinity brand 3
netperf 153
network 2, 120
network card 3
network drivers 127
network setup 29
NFS 257, 259
  access data 260
  COAS 257
  mounting a volume 110
  remote access 262
NIS 247
  additional information 256
  client 251
  COAS 249
  installation 247, 248
  server 249
number of chunks 58

O
odd parity 374
OpenLinux 2
open-source 1
orthogonal 384

P
packet filtering with IP chains 265
  additional information 278
  dail-up Internet connection 265

FTP masquerading 272
gateway 265, 273
checksum 273
demasquerade 274
forward chain 274
input chain 274
lo interface 274
local process 274
output chain 274
routing decision 274
sanity 274
IP chains 275
IP forwarding 268
NAT 265
network configuration 266
requirements 266
partition size 18
partitions
/boot 16, 17
DOS/Windows 16
extended 16, 18
logical drive 19
format 22
home 16
Linux 16
logical partitions 20
mounting points 17
opt 16
reset 20
root 16
sizes 387
swap 16
usr 16
var 16
performance
of RAID subsystems 380
performance tools in Linux 141
peripherals 114
Personal Systems Reference 3
personal systems reference (PSREF) 3
PHP 199
POP3 215, 234
printer 116
printer attributes 118
pstree 145

R

RAID 383
described 369
level 0 (RAID-0) 370
level 1 (RAID-1, RAID-1E) 371
level 10 (RAID-10) 372
level 5 (RAID-5) 373
level 5 enhanced (RAID-5E) 377
orthogonal RAID-5 379
performance 380
RAID Advisory Board 384
recommendations 382
software-based 369
summary of RAID levels 383
support for two disk failures 377
supported disk technologies 369
RAID level 58, 384
RAID performance characteristics 384
random access 297
random access memory 3
raw device backup 306
recompile kernel 4, 143
recovery 297
Red Hat 7.0 1
Red Hat Package Manager (RPM) 52
remote tape drive support 306
RFC 240
root password 26
RPC 258
RPM 36, 52, 98, 189, 198, 203, 241, 257, 288, 343
package management 97

S

Samba 155
COAS 155
configuration 157
global settings 158
installation 155
kpackage 156
NetBIOS 158
printer shares 165
shares 162
start 165
stop 166
SWAT 167
logon 169
WINS 155
Samba File System (SMBFS) 155
scheduling utility 304
SCSI adapter 3

SCSI drivers 128
SCSI host adapter drivers 128
Secure Shell (SSH) 279
  configuration 280
  host key generation 281
  installation 279
  sshd server daemon 281
  user key generation 281
Sendmail 213
  additional information 237
  configuration 226
  DNS configuration 216
  mail client 234
  mail routing 228
  MTA 213
  network configuration 215
  packages 213
  POP3 215, 234
  SMTP 215
  Webmin 229
server message block (SMB) 155
ServeRAID 4, 35, 36, 37, 55, 64, 75
  configuration 37
  driver 38
  firmware 4, 37
  hot swap rebuild 61
  ipssend 54
    commands 55
    devinfo 60
    getconfig 55
    getstatus 59
    hsrebuild 61
    rebuild 66
    setstate 63
    synch 65
    unattended 65
  rebuild drive 66
  remote management 78
  replace drive 67
  RPM 38, 39, 42, 43, 52
  synchronize logical drives 65
  unattended mode 65
  utility 38
ServeRAID adapter 4 37, 45
Simple Network Management Protocol (SNMP) 200, 285
  commands 285
  community strings 286
  network management application 285
  network performance 288
  object identifiers (OIDs) 286, 287
  SNMP agent 285
SMTP 215
sound drivers 128
source code 1
SSL 200
stripe unit size 58
support 2
SuSE 1
SWAT 167
  global settings 169, 171
  logon 169
  printers 179
  restart Samba 178
  Samba passwords 184
  Samba status 182
  shares 173
system load average 112

T

target partition 30
Tetris 32
text-based interface 91
The Open Source Development Lab 1
time 113
time zone 31
token-ring 4
  DHCP 85
  ibmtr 84
  ISA 84
  lanaid 85
  olympic 84
  PCI 84
token-ring drivers 128
top 144
TurboLinux 1

U

UNIX 1
user identification
  UID 102

V

video mode 385
  Netfinity 3000 385
  Netfinity 3500M10 385

Netfinity 5000 385
Netfinity 5500 M10 385
Netfinity 5500 M20 385
Netfinity 5600 385
Netfinity 7000 M10 385
Netfinity 8500R 385
video setup 10
virtual file support 306
vmstat 153

W

Web site
RAID Advisory Board 384
Webmin 129, 131, 220, 223, 226
actions log 132
configuration 132
help 132
IP forwarding 268
main window 131
sendmail 229
servers 132
users 132
write cache status 58

X

X.500 239, 240
XFree server 10, 11
xSeries brand 2, 3
X-Windows 78, 129

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 **Redbook Title**	SG24-5861-01 Caldera OpenLinux Integration Guide for IBM @server xSeries and Netfinity
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	O Very Good O Good O Average O Poor
Please identify yourself as belonging to one of the following groups:	O Customer O Business Partner O Solution Developer O IBM, Lotus or Tivoli Employee O 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.	O 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/

Caldera OpenLinux Integration Guide for IBM @server xSeries and Netfinity

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Caldera OpenLinux Integration Guide for IBM @server xSeries and Netfinity

The complete guide to running Caldera OpenLinux 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, sendmail, DNS, DHCP, LDAP, and more

Here's all the information you need to maximize Caldera OpenLinux performance and reliability on state-of-the-art IBM @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 OpenLinux 2.3/eServer deployment and system administration throughout the entire system life cycle!

You'll get running fast with expert step-by-step preparation and installation techniques: review updating your BIOS and firmware, making the CD-ROM bootable, preparing SCSI devices, partitioning, configuration, X-Windows setup, deploying IBM ServeRAID in OpenLinux environments, and much more.

Next, you'll master all the key techniques of day-to-day OpenLinux system administration, including backup and recovery, Internet and e-mail connectivity, DNS/DHCP name services, and using Caldera OpenLinux 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 OpenLinux on Netfinity servers.

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

SG24-5861-01 ISBN 0738419842