



SuSE Linux

Firewall on CD2

1st edition 2002

Copyright ©

This publication is intellectual property of SuSE Linux AG.

Its contents can be duplicated, either in part or in whole, provided that a copyright label is visibly located on each copy.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SuSE Linux AG, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

Many of the software and hardware descriptions cited in this book are registered trademarks. All trade names are subject to copyright restrictions and may be registered trade marks. SuSE Linux AG essentially adheres to the manufacturer's spelling. Names of products and trademarks appearing in this book (with or without specific notation) are likewise subject to trademark and trade protection laws and may thus fall under copyright restrictions.

Please direct suggestions and comments to documentation@suse.de

Authors: Michael Calmer, Uwe Gansert, Edith Parzefall, Ulrich Schairer,
Marius Tomaschewski, Paul Zirnik, Berthold Gunreben,
Karine Nguyen

Translators: Steve Tomlin

Editors: Antje Faber, Berthold Gunreben, Jana Jaeger, Edith Parzefall,
Peter Reinhart, Thomas Schraitle, Rebecca Walter

Layout: Manuela Piotrowski, Thomas Schraitle

Setting: L^AT_EX

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

Contents

Preface	1
1 Introduction	3
SuSE Firewall on CD 2	5
System Requirements	6
Network Planning	10
Security Policy and Communication Analysis	10
Security Policy	11
Communication Analysis	11
Typical Firewall Setups	12
2 SuSE Adminhost for Firewall	15
Installation of the SuSE Adminhost for Firewall	16
Language Selection	16
Selecting the Mouse	17
Keyboard and Time Zone	17
Preparing the Hard Disk	18
Boot Manager Configuration	20
Setting the root Password	20
Confirming Settings and Starting the Installation	21
Preparing the Graphical Interface	21
Configuring the Network with YaST2	22
Manual Network Configuration	22
Configuration Files	23
Start-Up Scripts	29
The User fwadmin for the FAS	30
Upgrade to the VPN Edition	31

3 Firewall Administration System (FAS)	33
Logging in as fwadmin	34
Starting the Firewall Administration System	34
Using the FAS	34
Initial Login	34
Creating User Accounts	34
Creating a New Configuration	36
Example, Inc.	38
The Setup	39
The Headquarters in Nuremberg	39
The Branch in Frankfurt	40
The Branch in Munich	40
Network Policies	41
Configuring the Base Setup	41
The Example, Inc., Configuration	50
IP Filter and NAT	53
IP Forward	53
Masquerading	54
Destination NAT	54
ICMP to Firewall	55
Kernel Runtime Setup	59
System Logging	60
DNS Forwarder	61
FTP Proxy — External	63
FTP Proxy — Internal	66
Generic TCP Proxy	69
Configuring the HTTP Proxy — External	72
Configuring the HTTP Proxy for Internal Connections	74
IPsec VPN Tunnel	84
Configuring the Mail Relay	94
Administration via SSH	97
Time Synchronization	99

Log File Analysis	99
The Log Files	101
Evaluating the Log Files	102
The IP Filter Statistics	102
The Interface Statistics	104
Mail Statistics	104
Certificate Management	105
Creating a Certificate Authority	106
Creating a Certificate	107
Deleting a Certificate	107
Importing Certificates	107
Exporting Certificates	108
Saving the Configuration	109
Editing an Existing Firewall Configuration	109
Editing Configuration Files	109
Testing the Configuration	110
Documenting Configuration, Tests, and Results	111
Monitoring the Firewall	111
4 SuSE Live CD for Firewall	113
Hardware Requirements	114
Description	115
Services on the Firewall	115
iptables	117
FreeS/WAN	121
DNS	123
Mail	123
HTTP Proxy	123
FTP Proxy	124
SSH	125
chroot, compartment, Kernel Capabilities	125
The Configuration Disk	125
Creating the Configuration Disk	126
The Configuration Files	126
Boot Parameters	130

5	IPsec Client on Windows XP and Windows 2000	131
	Exporting the Required Certificates	131
	Importing the Certificates in Windows	131
	Configuring the Required Snap-Ins	132
	Importing the Client Certificate	132
	Making a Note of Important Certificate Data	133
	Configuring the IPsec Connection	133
	Installing the ipsec.exe Tool	133
	Editing ipsec.conf	134
	Creating a Desktop Link and Activating the Connection	134
	Closing the Connection	135
6	Implementing the Firewall	137
	Requirements for Successful Implementation	138
	Booting the Firewall Host	138
	Testing the Firewall	138
	Internal Testing	139
	External Testing	140
	Going Online	140
7	Help	141
	Troubleshooting	142
	Problems Installing the Adminhost	142
	Problems Booting the Live CD	142
	Problems Integrating the Network	142
	Detecting Attacks	143
	Intrusion Detection and Event Display	143
	External Attacks	146
	Advantage of the Live File System of the SuSE Firewall on CD	147
	Recommended Reading	148

8 Support, Maintenance, and Patch Management	149
Maintenance	149
Accessing the SuSE Maintenance Web	149
Getting Patches	149
Support and Services	152
Support Conditions	152
Commercial Support	153
SuSE Training Program	154
Feedback	155
Additional Services	155
A DNS — Domain Name Service	157
Starting the Name Server BIND	158
The Configuration File /etc/named.conf	159
Important Configuration Options	160
The Configuration Section “Logging”	161
Zone Entry Structure	162
Structure of Zone Files	163
For More Information	166
B Proxy Server: Squid	167
What is a Proxy Cache?	168
Some Facts About Cache Proxying	168
Squid and Security	168
Multiple Caches	169
Caching Internet Objects	169
System Requirements	170
Hard Disk	170
RAM	171
CPU	172
Starting Squid	172
The Configuration File /etc/squid.conf	173
Transparent Proxy Configuration	178

Kernel Configuration	179
Configuration Options in /etc/squid.conf	179
Squid and Other Programs	179
SquidGuard	179
Cache Report Generation with Calamaris	181
More Information on Squid	182
C Network Security	183
Masquerading and Firewalls	184
Masquerading Basics	184
Firewalling Basics	186
SuSEfirewall	187
SSH — Secure Shell, the Safe Alternative	190
The OpenSSH Package	190
The ssh Program	190
scp — Secure Copy	191
sftp — Secure File Transfer	191
The SSH Daemon (sshd) — Server-Side	192
SSH Authentication Mechanisms	193
X, Authentication, and Other Forwarding Mechanisms	194
Security and Confidentiality	195
Basic Considerations	195
Local Security and Network Security	195
Some General Security Tips and Tricks	205
Using the Central Security Reporting Address	208
D YaST and SuSE Linux License Terms	209
E The GNU General Public License	213

Preface

Many thanks to Jürgen Scheiderer, Carsten Höger, Remo Behn, Thomas Biege, Roman Drahtmüller, Marc Heuse, and Stephan Martin.

The SuSE Firewall on CD 2

The SuSE Firewall on CD is a tool package allowing set up of a firewall solution for your network. It helps with the related configuration, monitoring, and administration chores. The complete functionality of the SuSE Firewall on CD is based on Open Source programs specially selected and enhanced for this purpose.

The SuSE Firewall on CD protects your local network from illegal access, manages authorized use of your resources, and provides a controlled channel through which users on your local network are enabled to communicate with the Internet.

To reduce the likelihood of the firewall being misconfigured, we recommend using the Firewall Administration System (FAS), which corrects most user mistakes by performing a number of sanity checks. FAS ensures a consistent configuration without limiting the available options.

Total security cannot be achieved by any firewall, not even by the SuSE Firewall on CD. This package is not intended to serve as a substitute for a proper security policy and does not eliminate the need to administer, maintain, and monitor the firewall. It provides a number of tools to make these tasks much easier to perform.

Note

Although every effort has been made to make SuSE Firewall on CD a secure firewall product, SuSE is unable to guarantee the security of your system. SuSE cannot be held responsible for any damages caused by insufficient security measures.

Note

Introduction

The importance of the Internet, the communication possibilities provided, and the information-gathering options offered seem to grow every day. The number of companies and individuals with access to this worldwide network is growing steadily. However, connecting to the Internet often introduces some security concerns that should not be underestimated. The SuSE Firewall on CD can minimize and control the risks involved.

SuSE Firewall on CD 2	5
System Requirements	6
Security Policy and Communication Analysis	10
Typical Firewall Setups	12

Most companies rely on their own networks to exchange and process mission-critical information for in-house purposes, such as an intranet, databases, and e-mail. Without the proper protection mechanisms in place, all this data would be widely available to the outside world as soon as the local network was connected to the Internet — something that could obviously cause a lot of damage, especially for companies.

It's easy to run a secure computer system. You just have to disconnect all dial-up connections and permit only direct-wired terminals, put the machine and its terminals in a shielded room, and post a guard at the door.

F.T. Grampp and R.H. Morris

In the real world, it is impossible to run a computer system in this way. If you are an administrator, you will know all too well about the problems arising from the increasing number of networked machines. After all, you are the person who has to deal with the situation on a daily basis. Up to now, your network may have been quite manageable. It was once possible to know the users on a network and provide network functionality to your users mainly on the basis of trust relationships, which kept the administrative overhead at a reasonable level. Now, with connections to the Internet commonplace, things have changed dramatically, as have the duties of the administrator. All at once, there are users who are completely unknown who can use resources on your network, such as the web server.

These users need to be handled in a totally different manner than your company's internal employees. On the other hand, it has been proven that eighty percent of all attacks on corporate networks are not carried out from the Internet. There are other reasons to protect the corporate network and to restrict access to it: crackers, or intruders, are always on the lookout for places to store pirated software or, even worse, contents prohibited by law. Not taking any measures against this could not only mean that the company could open itself up to legal prosecution, but could also put your company's good reputation at stake.

Today, there are a number of different mechanisms available to prevent unauthorized access to corporate networks and resources (i. e., disk space or CPU capacity), from IP packet filters on routers to multiple-level firewall solutions complete with a demilitarized zone (DMZ). In a general sense, of course, the word "firewall" is used for a device prevent the spreading of fires. Firewalls made of bricks are found in buildings where they are used to isolate whole sections from each other and in cars a special metal plate shields the passenger compartment from the engine compartment. Similarly, the purpose of Internet

firewalls is to fend off attacks directed at your intranet as well as to regulate and protect clients on your LAN by imposing an access policy.

The first firewall was a non-routing UNIX host connected to two different networks: one network interface was connected to the Internet and the other one to a private LAN. To reach the Internet from within the private network, users had to log in to the UNIX firewall server before they could access any outside host. To do so, they would start, for example, an X Window–based browser on the firewall host then export the window to the display of their workstation. With the browser hosted on the firewall, users had access to both systems at the same time. However, you should not consider this kind of setup (called “dual homed system”) for your own network unless you really trust all the users on it. To understand why, remember that ninety-nine percent of all break-ins on computer systems start with an attempt to obtain a user account on the targeted system.

The scope of your firewall solution depends on the required degree of protection, which may need to be in line with legal regulations in some cases and which should be determined through a communication analysis. Information about conducting a communication analysis yourself or obtaining consulting, training, or support services is available in [Help](#) on page 141.

Another factor affecting the operation of a firewall is the state of its documentation. There should be a way to determine “who has changed what, when, and how” to trace back whether changes were made by an authorized party. This will also be quite helpful when it comes to dealing with things like certifications and audits.

The SuSE Firewall on CD is a product covering the whole range of these issues in all its details, from packet filtering to setting up a multiple-level firewall. Because the package is based on Open Source programs, it is also possible to audit the source code without too much difficulty.

SuSE Firewall on CD 2

The SuSE Firewall on CD consists of the following two product components:

1. the SuSE Live CD for Firewall
2. the SuSE Admin CD for Firewall

For the sake of simplicity, we mostly refer to these in this manual as “Live CD” and “Admin CD”. The Live CD contains the firewall package proper, which, in

our case, is based on the concept of an application-level gateway combined with IP packet filtering. The firewall's routing and gateway capabilities are turned off by default, but can be enabled as required. All requests accepted and processed by the firewall are handled at the application level. The package supports the most important Internet protocols: SMTP, FTP, HTTP, HTTPS, DNS, and VPN.

With the Admin CD, install the SuSE Adminhost, which takes care of the configuration, monitoring, and administration of the SuSE Firewall on CD. To properly manage these tasks, the CD includes a special selection of software packages and the necessary installation routines.

Connecting an internal company network to the Internet requires a good deal of planning, including a security concept based on a communication analysis and a demand analysis. Furthermore, there should be a concept describing how to deal with emergencies, such as break-ins or data loss. Also make sure the firewall machine is protected from unauthorized physical access. The best way to achieve that is to lock the machine in a separate server room.

System Requirements

Motherboards: all except EISA, VLB

Processors: from Pentium I and compatible (e.g., AMD, Cyrix, IBM)

Main memory:

- for the Admin CD: at least 128 MB RAM
- for the Live CD: 256 MB RAM

Hard drive:

- for Admin CD: min. 2 GB
- for the Live CD: none or a large one for log files (10 GB recommended)

CD-ROM:

- IDE drive: all
- SCSI drives: all CD-ROM drives connected to one of the SCSI host adapters listed below

IDE Controller: all except IDE-RAID controller

ISDN Cards: ▪ AVM Fritz!PCI v2.0

- ELSA Quickstep 1000 PCI
- Generic HFC 2BDSO PCI

SCSI Host Adapters:

53c7,8xx: NCR 53c7,8xx (old driver)

AM53C974: AM53/79C974

BusLogic: BusLogic

DAC960: Mylex DAC-960/DAC1100

a100u2w: Initio INI-A100U2W

aacraid: Adaptec RAID

advansys: AdvanSys

aha152x: Adaptec 1505/151x/152x/2825

aha1542: Adaptec 154x

aha1740: Adaptec 1740

aic7xxx: Adaptec 274x/284x/294x

atp870u: ACARD AEC-671X

cciss: Compaq CISS Array

cpqarray: Compaq SMART2 RAID

dc395x_trm: Tekram Tekram DC395U/UW/F, DC315/U

dpt_i2o: DPT I2O SmartRAID V

dtc: DTC 3180/3280

eata: EATA (E)ISA (PM2011/021/012 etc)

eata_pio: EATA-PIO (old DPT PM2001, PM2012A)

fdomain: Future Domain 16xx

gdth: ICP Vortex GDTH Disk Array

in2000: Always IN 2000

initio: Initio INI-9X00U/UW

ips: IBM ServeRAID

megaraid: AMI Megaraid

ncr53c8xx: NCR 53c8xx

pci2000: PCI-2000

pci2220i: PCI-2220I

psi240i: PSI-240i
qlogicfas: Qlogic FAS
qlogicfc: QLogic ISP 2100 SCSI-FCP
qlogicisp: QLogic ISP 1020
qlogicpti: PTI Qlogic ISP Driver
seagate: Seagate ST-02/Fut. Domain TMC-8xx
sim710: Simple 53c710 (Compaq, NCR)
sym53c416: Symbios 53C416
sym53c8xx: Symbios 53c8xx
t128: Trantor T128/T128F/T228
tmscsim: Tekram DC390(T) (AM53C974 chip)
u14-34f: UltraStor 14F/34F
ultrastor: UltraStor (alternate driver)
wd7000: Western Digital 7000 FASST

Network drivers:

3c501: 3Com 3c501
3c503: 3Com 3c503
3c505: 3Com 3c505
3c507: 3Com 3c507
3c509: 3Com 3c509/3c579
3c515: 3Com 3c515
3c59x: 3Com 3c59x/3c90x (592/595/597)
3c90x: 3Com 3c90x/3c980 B/C series
82596: i82596 Ethernet Driver
acenic: Alteon AceNIC/3C985/NetGear GA620
arlan: Aironet Arlan 655
at1700: AT1700
cs89x0: CS89x0
de4x5: DE425, DE434, DE435, DE450, DE500
de600: D-Link DE600 pocket adaptor
de620: D-Link DE620 pocket adaptor

depca: DEPCA,DE10x,DE200,DE201,DE202,DE422
dgrs: Digi Intl. RightSwitch SE-X
dmfe: DM9102 PCI Fast Ethernet
e100.o: EtherExpress PRO/100 (Intel driver)
e1000.o: Intel(R) PRO/1000 Gigabit Server Adapter
e2100: Cabletron E21xx
eeepro100: Intel EtherExpress Pro 100
eeepro: Intel EtherExpressPro
eexpress: Intel EtherExpress 16
epic100: SMC 83c170 EPIC/100
eth16i: ICL EtherTeam 16i/32
ewrk3: EtherWORKS 3 (DE203, DE204, DE205)
fmv18x: FMV-181/182/183/184
hamachi: Packet Engines GNIC-II PCI Gigabit Ethernet Adapter
hp-plus: HP PCLAN+ (27247B and 27252A)
hp100: HP 10/100VG PCLAN (ISA, PCI)
hp: HP PCLAN (27245 / 27xxx)
lance: AMD LANCE and PCnet (AT1500/NE2100)
ne2k-pci: NE 2000 (PCI)
ne: NE 2000 / NE 1000 (ISA)
ni5010: NI5010
ni52: NI5210
ni65: NI6510 (am7990 lance chip)
old_tulip: DEC Tulip (DC21x4x) (alter Treiber)
pcnet32: AMD PCI PCnet32 (PCI bus NE2100)
r1100a: Compex RL-100ATX
rtl8139: RealTek RTL8129/8139
sis900: SiS 900 PCI Fast Ethernet
sk98lin: SysKonnnect Gigabit Ethernet 984x
skfp: SysKonnnect FDDI (SK-55xx/SK-58xx)
smc-ultra32: SMC Ultra32
smc-ultra: SMC Ultra

smc9194: SMC 9194
tlan: Compaq Netelligent 10/100/NetFlex 3
tulip: DEC Tulip (DC21x4x) PCI
via-rhine: VIA VT86c100A Rhine-II
wd: Western Digital WD80x3
yellowfin: Packet Engines Yellowfin Gigabit

Network Planning

Before beginning the installation of the Adminhost and the configuration of the firewall, consider your network layout. The diagrams in the following section provide some ideas for layout options. The most important thing to do before you begin installing the software is to consider what hardware to use for each part of the system.

The actual firewall host is a very special server. SuSE Firewall on CD should not be used on your main server or another computer in the network. It is best if the firewall host does not have a hard disk, although one is required for using Squid and similar programs. It needs, however, a floppy disk drive for the configuration disk, a bootable CD-ROM drive for booting the Live CD, and network interfaces. The firewall host can serve as a gateway for your system.

The Adminhost should be a dedicated machine. It is used to create the configuration floppy for the firewall host. It should be possible to run a graphical user interface on this machine, so you can use the Firewall Administration System (FAS) to configure the firewall and create the configuration disk. All needed programs are included on the Admin CD. It can also be used as the log host.

The log host is a machine used to log the events on the firewall. It needs a large hard disk for storing information. This does not have to be a dedicated machine. It is, however, recommended to make it a dedicated machine, for reasons of security, and also use it as the Adminhost. It should never be unavailable to the firewall host.

Security Policy and Communication Analysis

To ensure that the internal network's connection to the Internet (or to any other "unprotected" network) is secure, a few things need to be clarified first. This includes outlining a security policy for your own network and undertaking a communication analysis.

Security Policy

The security policy provides the basis for working with all programs, hosts, and data. In addition, it outlines how to guarantee the monitoring of security guidelines and how internal or external breaches of this policy are handled. To draft a security policy, it is best to produce a communication analysis. To this end, the following topics are of utmost importance:

- An analysis of your security requirements is necessary. What needs to be protected?
- Are there areas of the intranet that contain especially sensitive data (such as the personnel department and data critical of the company)? Where is this data located?
- Who can access the data? Are there various levels of authorization?
- Should data be available over the network?
- Which services should be accessible internally? Which services should be accessible from internal to external (e-mail, surfing, data transfer) and which services external to internal (e-mail, web services, data transfer)?

The list of questions to answer in a security policy must be drawn up individually and answered.

Communication Analysis

The most important aid for carrying out a communication analysis is a communication matrix. The services available for client hosts and users are represented here in a table format. This matrix is then mapped to the proxies and IP filter rules.

Setting up a Communication Matrix

Make a list of all clients and servers on your network. Then define which protocols may be used by which clients. Also state in which direction each packet can be sent or received.

An example for the HTTP protocol: The client host1 needs to access a web server in the internal network, but should not be able to establish a connection to an external web server. The entry in the communication matrix then appears as depicted in the example shown below.

Protocol	icmp		ftp		ssh		smtp		http		https		...
Client	internal	external	i.	e.	i.	e.	i.	e.	i.	e.	i.	e.	
host1	X	-	X	-	-	-	X	-	X	-	-	-	
host2	X	-	-	-	X	-	X	-	-	-	-	-	
host3													
...													
hostn													

With the help of such a communication matrix, obtain an overview of the communication constellations within the network. This simplifies the configuration of your network and error analysis.

Typical Firewall Setups

This section gives a brief overview of the most typical firewall setups. All the configurations presented below can be implemented with the SuSE Firewall on CD.

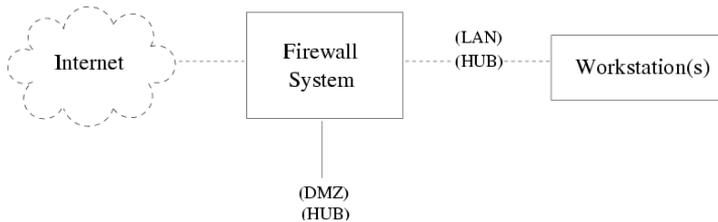


Figure 1.1: Very Basic Setup

Figure 1.1 shows a firewall with three network interfaces: an external one to connect to the Internet and two internal ones connecting with the corporate network via LAN or HUB and with the DMZ (demilitarized zone) using another HUB. With this setup, the firewall has to perform all the functions of the default gateway (router) and the packet filter. The internal network would be left completely open if the firewall were infiltrated.

The setup shown in Figure 1.2 on the facing page is still a relatively simple one. The DMZ is only protected by a packet filter on the router (screening router).

The setup shown in Figure 1.3 on the next page is still not very complicated, but provides much better protection than the previous ones. A “screening router”

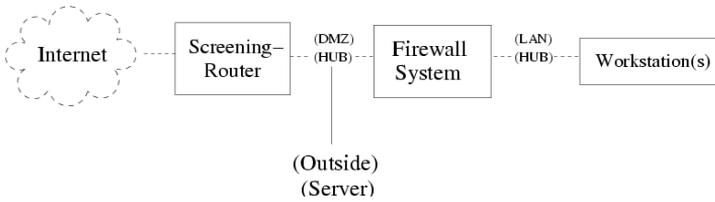


Figure 1.2: Simple Setup

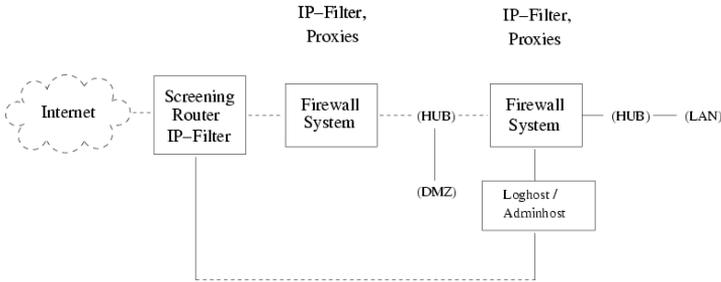


Figure 1.3: Effective and Manageable Setup

stops any illegal requests at the packet level. Packets allowed through are passed to the first firewall host, which operates at the application level and uses packet filtering rules to control access to the DMZ, to the second firewall host, and to the internal network beyond. This second firewall is a packet filtering proxy machine protecting the internal network. In addition to this, it controls access from the internal network to the Internet and the DMZ.

SuSE Adminhost for Firewall

It is no easy task to administer, maintain, and monitor a firewall. Above all, the importance of monitoring the firewall should not be underestimated. This is why the SuSE Firewall on CD includes the SuSE Adminhost for Firewall, which helps with configuring, administering, and maintaining the firewall.

Installation of the SuSE Adminhost for Firewall	16
Configuring the Network with YaST2	22
Manual Network Configuration	22
The User <code>fwadmin</code> for the FAS	30
Upgrade to the VPN Edition	31

After installing the SuSE Adminhost for Firewall, the Firewall Administration System (FAS) is available. The FAS is a tool with a graphical administration interface allowing menu-driven configuration of the SuSE Firewall on CD.

Configurable tools are available for monitoring the firewall that can perform tests of the firewall, evaluate the log files, and monitor network traffic. In the case of an attack, there is the possibility of informing system administration by e-mail, pager, or other means, so suitable counter measures can be taken as quickly as possible (for example, cutting off the network connection to the Internet or intranet).

A firewall without any monitoring cannot provide any real protection. Also, all changes made to the firewall configuration should always be thoroughly documented. By means of this documentation, errors can be found more easily or modifications made by unauthorized parties can be undone.

If you are already running a SuSE Firewall on CD (Standard or VPN edition) and administer it with SuSE Adminhost for Firewall, you must still perform a new installation of the Adminhost, because version 2 runs on the 2.4 kernel. This means that the packet filter `iptables` is now used instead of `ipchains`. Unfortunately, as a consequence, any configuration used previously can *no longer* be used. If you have, in addition to the SuSE Firewall on CD, purchased the VPN module or if you may want to install this at a later date, read the instructions in [Upgrade to the VPN Edition](#) on page 31.

Installation of the SuSE Adminhost for Firewall

To install the Adminhost for the first time, insert this CD into the CD drive of the computer then reboot the machine. If the computer does not boot from the CD, change the boot sequence in the BIOS accordingly.

The `linuxrc` welcome screen appears. Simply press  and the automatic hardware detection starts. `YAST2` leads you through the rest of the installation. All the necessary programs for administration, configuration, and maintenance are installed.

Language Selection

The first decision to make in the installation process is the language. Either use the mouse or the keyboard. All entry fields, selection lists, buttons, and switches can be selected by clicking with the mouse. To use the keyboard, the following rules apply:

- **(Tab)** moves the focus forward an entry or selection field or a button. **(Shift) + (Tab)** moves the focus to the previous item. With **(↑)** and **(↓)**, depending on which area is activated, make a selection or cycle through a list.
- With **(↵)**, the selected command is carried out — the action shown on the active button.
- With **(Space)**, entries can be marked.
- In addition, most actions can be started with the key combination **(Alt) + the underlined letter**.

Tip

Here and in the following dialogs, YaST2 is just collecting information. Later, YaST2 displays the information it has collected. In 2 on page 21, you still have the chance, by means of the 'Back' button, to return to the previous dialogs and correct details.

Tip

Select your preferred language. When you have chosen a language, select 'Apply' to switch all texts to the selected language.

Selecting the Mouse

This dialog only appears if YaST2 was not able to detect the mouse automatically. A dialog window with a long list of mouse types appears from which to select the appropriate mouse type.

Once you have found the right mouse type, move with **(Tab)** to 'Test' and press **(↵)**. Now move the mouse. If the mouse cursor moves normally, everything is working properly and you can click 'Next'. If the mouse does not work, return to the selection list using **(Tab)** and change the settings.

If no mouse type functions or if you do not want to use a mouse, activate 'No mouse'. The rest of the installation is carried out using only the keyboard.

Keyboard and Time Zone

Select the keyboard layout and time zone.

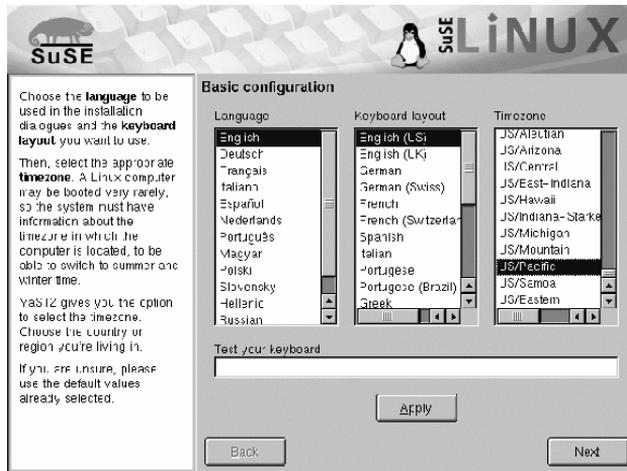


Figure 2.1: YaST2: Keyboard Layout and Time Zone

- Now test your keyboard. By clicking with the mouse or using **(Tab)**, activate the entry line and type in letters there. Especially test ‘y’, ‘z’, and special characters.
- The second item is a list of countries in a tree structure (continent/-country/region). Select your country or region from these. YaST2 finds the appropriate time zone.

Use ‘Next’ to proceed to the next dialog window.

Preparing the Hard Disk

In the following steps, select the hard disk or disks and prepare them for the installation. Depending on your computer’s hardware, there may be slight differences from the dialogs appearing here.

Step 1

If more than one hard disk exists, decide which one to use for the installation. The disks found will be listed. See Figure 2.2 on the facing page. Select the last option (‘Custom Partitioning’) to partition them by hand. In this dialog, it is also possible to continue using existing partitions by specifying formatting and allocating mount points directly. Normally, choose *one* hard disk then click ‘Next’.

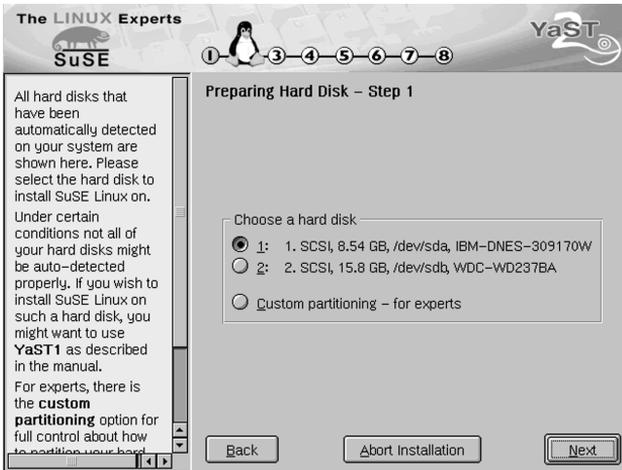


Figure 2.2: YaST2: Preparing the Hard Disk

Step 2

One of the following situations could occur:

- If the hard disk is *not* empty, YaST2 shows all existing partitions on the hard disk as well as ‘Use entire hard disk’. *Free, unpartitioned* storage space at the “end” of the hard disk is also displayed and is automatically preselected. YaST2 can use further space for SuSE Linux, but only if it is contiguous — partitions can only be released for further use “from behind”. For example, if three partitions exist, partitions 1 and 2 remain and you must select partition 3. To make the entire hard disk available for SuSE Linux, select ‘Entire Hard Disk’.
- For an *empty* hard disk, the entire hard disk is used for SuSE Linux.

If you have other requirements, press ‘Back’ to return to the previous dialog, as mentioned on the preceding page, for manual partitioning with the help of ‘Extended Settings’.

Note

Because the partitions made available for SuSE Linux will be formatted, all existing data there will be irretrievably lost.

Note

Once the installation starts and all requirements have been fulfilled, YaST2 partitions and formats the necessary hard disk space on its own. The entire hard disk or the available partitions are split up for SuSE Linux into the three standard partitions: a small partition for /boot (about 16 MB) as close as possible to the beginning of the hard disk, a partition for swap (128 MB), and all the rest for / (root partition).

Boot Manager Configuration

For Linux to be bootable after the installation, a boot mechanism must be installed. Specify how the boot manager LILO "Linux LOader" should be installed or if another boot procedure should be used.

Setting the root Password

The user `root` has special privileges in Linux. He can, for instance, start and stop system processes, create and remove users, and manipulate important system files — in other words, perform the duties of the system administrator.

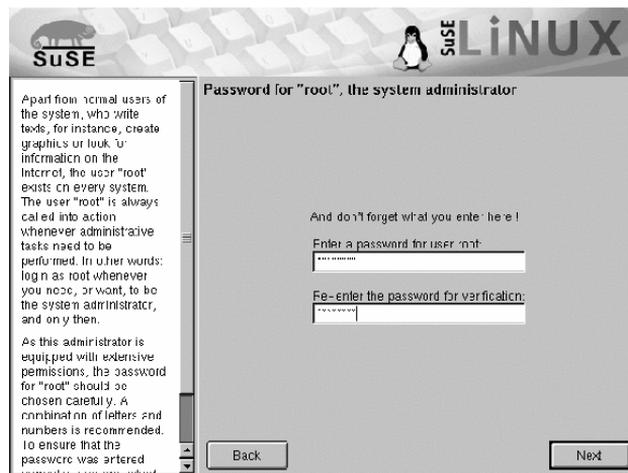


Figure 2.3: YaST2: Entering the root Password

Provide a password for the user `root`.

Note

Remember the `root` password very carefully, as you cannot retrieve it later. This password whenever you perform administrative tasks on the system.

Note

If you press 'Next', the installation will start.

Confirming Settings and Starting the Installation

Review all the settings made until now. To make changes, cycle through the windows with 'Back'. If you press 'Next', you are asked again for confirmation (in green) to start the installation with the settings as shown. After confirming with 'Yes — install', YaST2 begins setting up the system. With 'No', you have the option of checking data again and changing items where necessary by pressing 'Back' to reach the relevant window.

You still have the option of aborting the installation completely. All settings made and details supplied will be lost. If you select 'Abort installation', your computer, after asking for confirmation, shuts down and you can switch off the computer or reboot without any problem. Up to this point, no changes have been made to your computer.

With 'Save Settings to Floppy Disk', save all details to floppy disk and use them again for other installations. This can only be selected if it is supported by your hardware.

After selecting 'Yes — install', partitions are created and formatted. Depending on your system's capacity and the size of the hard disk, this could take some time. Avoid aborting here, as this would put the hard disk into an undefined condition.

Afterwards, the packages from the CD are loaded and the SuSE Linux base system is installed. After you have confirmed this with 'OK', the base system is started. YaST2 continues the installation of software and, if necessary, requests additional CDs. If you 'Abort' the installation at this stage, the system is unusable.

Configure the graphical interface then try out SuSE Linux.

Preparing the Graphical Interface

To provide a graphical user interface, YaST2 tries to find the information it needs for the monitor and graphics card. If this is successful, a sensible screen

resolution, color resolution, and repetition rate frequency are selected for the monitor and a test screen is displayed.

Note

Check the settings before accepting. If you are not sure, consult the documentation for your graphics card and monitor.

Note

If the monitor is not detected, select your model from the list provided. If you have an unknown model, enter the settings by hand or have the data loaded from a “driver disk” provided with your monitor. Consult the documentation for your monitor.

Set up the required screen resolution and a color depth of 16 bpp. Check the settings by choosing ‘Test’ and make adjustments if necessary.

Tip

In rare cases, it may be necessary to configure the X server “by hand.” To do this, run the program `SxX` at a later time.

Tip

The default window manager on the Adminhost is KDE2.

Configuring the Network with YaST2

Have the following data on hand when configuring the network: IP address, network mask, and default gateway. If you are managing a DHCP server, you can also configure the Adminhost as a DHCP client. However, you *must* assign a static IP address to the Adminhost.

Select the network card from those detected by the hardware recognition. Enter the IP address and network mask. Enter all the required information if a name server already exists, as shown in Figure 2.5 on page 24.

Manual Network Configuration

Manual configuration of the network software is a last resort. We recommend using YaST, but YaST cannot cover all aspects of network configuration, which requires, in some cases, a bit of manual fine-tuning.

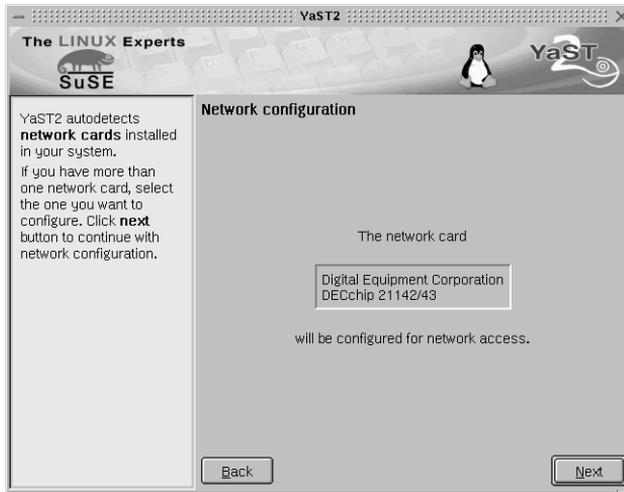


Figure 2.4: YaST2: Network Configuration

Configuration Files

This section provides an overview of the network configuration files. It explains their purpose as well as the format used.

`/etc/rc.config`

The majority of the network configuration takes place in this central configuration file. When making changes via YaST or when running SuSEconfig after the file has been modified manually, most of the following files are automatically generated from these entries. Even the boot scripts are configured via these file entries.

Tip

If you modify this file by hand, run SuSEconfig separately afterwards so the modified configurations are entered into the correct files.

Tip

`/etc/hosts`

In this file (see File 1 on the next page), IP addresses are assigned to host names. If no name server is implemented, all hosts to which an IP connection will be made must be listed here. For each host, a line

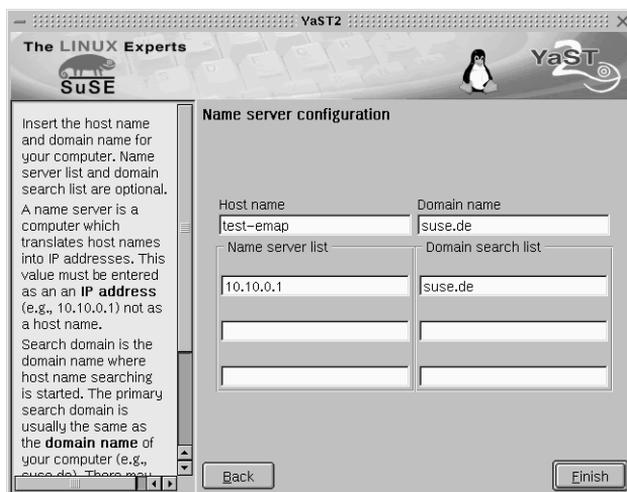


Figure 2.5: YaST2: Configuring the Name Server

consisting of the IP address, the fully qualified host name, and the host name (e.g., earth) is entered into the file. The IP address must be at the beginning of the line. The entries should be separated by blanks and tabs. Comments are always preceded by the '#' sign.

```
#
# hosts          This file describes a number of host name-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time when no name servers are running.
#               On small systems, this file can be used instead of a
#               "named" name server.  Just add the names, addresses,
#               and any aliases to this file.
#
127.0.0.1 localhost
192.168.0.1 sun.cosmos.com sun
192.168.0.20 earth.cosmos.com earth
# End of hosts
```

File 1: /etc/hosts

/etc/networks

Here, network names are converted to network addresses. The format is similar to that of the `hosts` file, except the network names precede the addresses (see File 2).

```
#
# networks This file describes a number of net name-to-address
# mappings for the TCP/IP subsystem. It is mostly
# used at boot time, when no name servers are running.
#
loopback 127.0.0.0
localnet 192.168.0.0
# End of networks.
```

File 2: /etc/networks

/etc/host.conf

Name resolution — the translation of host and network names via the *resolver* library — is controlled by this file. This file is only used for programs linked to the libc4 or the libc5. For current glibc programs, refer to the settings in */etc/nsswitch.conf*. A parameter must always stand alone in its own line and comments preceded by a '#' sign. Table 2.1 shows the parameters available.

<code>order <i>hosts, bind</i></code>	Specifies in which order the services are accessed for name resolution. Available arguments are (separated by blank spaces or commas): <i>hosts</i> : Searches the <i>/etc/hosts</i> file <i>bind</i> : Accesses a name server <i>nis</i> : Via NIS
<code>multi <i>on/off</i></code>	Defines if a host entered in <i>/etc/hosts</i> can have multiple IP addresses.
<code>nospoof <i>on</i></code> <code>alert <i>on/off</i></code>	These parameters influence the name server <i>spoofing</i> , but, apart from that, do not exert any influence on the network configuration.
<code>trim <i><domainname></i></code>	The specified domain name is separated from the host name following the host name resolution (as long as the host name includes the domain name). This option is useful if only names from the local domain are in the <i>/etc/hosts</i> file, but should still be recognized with the attached domain names.

Table 2.1: Parameters for */etc/host.conf*

An example for `/etc/host.conf` is shown in [File 3](#).

```
#
# /etc/host.conf
#
# We have named running
order hosts bind
# Allow multiple addrs
multi on
# End of host.conf
```

File 3: /etc/host.conf

/etc/nsswitch.conf

With the GNU C Library 2.0, the “Name Service Switch” (NSS) became more important. See the man page for `nsswitch.conf` or, for more details, *The GNU C Library Reference Manual*, Chapter “System Databases and Name Service Switch”. Refer to package `libcinfo`, series `doc`.

In the `/etc/nsswitch.conf` file, the order of certain data is defined. An example of `nsswitch.conf` is shown in [File 4](#). Comments are preceded by ‘`#`’ signs. Here, for instance, the entry under “database” `hosts` means that a request is sent to `/etc/hosts` (files) via DNS (see [A](#) on page [157](#)).

```
#
# /etc/nsswitch.conf
#
passwd:      compat
group:       compat

hosts:       files dns
networks:    files dns

services:    db files
protocols:   db files

netgroup:    files
```

File 4: /etc/nsswitch.conf

The “databases” available over NSS are listed in Table 2.2. In addition, `automount`, `bootparams`, `netmasks`, and `publickey` are expected in the near future.

<code>aliases</code>	Mail aliases implemented by <code>sendmail</code> ; see also the man page for <code>aliases</code> .
<code>ethers</code>	Ethernet addresses.
<code>group</code>	For user groups, used by <code>getgrent</code> ; see also the man page for <code>group</code> .
<code>hosts</code>	For host names and IP addresses, used by <code>gethostbyname</code> and similar functions.
<code>netgroup</code>	Valid host and user lists in the network for the purpose of controlling access permissions; see also the man page for <code>netgroup</code> .
<code>networks</code>	Network names and addresses, used by <code>getnetent</code> .
<code>passwd</code>	User passwords, used by <code>getpwent</code> ; see also the man page for <code>passwd</code> .
<code>protocols</code>	Network protocols, used by <code>getprotoent</code> ; see also the man page for <code>protocols</code> .
<code>rpc</code>	“Remote Procedure Call” names and addresses, used by <code>getrpcbyname</code> and similar functions.
<code>services</code>	Network services, used by <code>getservent</code> .
<code>shadow</code>	“Shadow” user passwords, used by <code>getspnam</code> ; see also the man page for <code>shadow</code> .

Table 2.2: Available “Databases” via `/etc/nsswitch.conf`

The configuration options for NSS “databases” are listed in Table 2.3 on the following page.

files	directly access files, for example, to <code>/etc/aliases</code> .
db	access via a database.
nis	NIS
nisplus	
dns	Only usable by <code>hosts</code> and <code>networks</code> as an extension.
compat	Only usable by <code>passwd</code> , <code>shadow</code> , and <code>group</code> as an extension.

Table 2.3: Configuration Options for NSS “Databases”

also it is possible to trigger various reactions with certain lookup results; details can be found in the man page for `nsswitch.conf`.

`/etc/nscd.conf`

The `nscd` (Name Service Cache Daemon) is configured in this file (see the man pages for `nscd` and `nscd.conf`). This affects the data resulting from `passwd`, `groups`, and `hosts`. The daemon must be restarted every time the name resolution (DNS) is changed by modifying the `/etc/resolv.conf` file. The command `rcnscd restart` serves this purpose.

Caution

If, for example, caching for `passwd` is activated, it usually takes about fifteen seconds until a newly added user is recognized by the system. By resarting `nscd`, reduce this waiting period.

Caution

`/etc/resolv.conf`

As is already the case with the `/etc/host.conf` file, this file, by way of the *resolver* library, plays a role in host name resolution.

Specified are the domain to which the host belongs (keyword `search`) and the status of the name server address (keyword `name server`) to access. Multiple domain names can be specified. When resolving a name that is not fully qualified, an attempt is made to generate one by attaching the individual search entries. Multiple name servers used by entering several lines, each beginning with `name server`. Comments are preceded by `#` signs.

An example `/etc/resolv.conf` is shown in File 5.

```
# /etc/resolv.conf
#
# Our domain
search cosmos.com
#
# We use sun (192.168.0.1) as name server
name server 192.168.0.1
# End of resolv.conf
```

File 5: /etc/resolv.conf

YaST enters the given name server here.

/etc/HOSTNAME

Here is the host name without the domain name attached. This file is read by several scripts while the machine is booting. It may only contain one line where the host name is mentioned. This file will also automatically be generated from the configuration in `/etc/rc.config`.

Start-Up Scripts

Apart from the configuration files described above, there are also various scripts that load the network programs while the machine is booted. This is started as soon as the system is switched to one of the *multiuser runlevels* (see also Table 2.4 on the following page).

<code>/etc/init.d/network</code>	This script takes over the configuration for the network hardware and software during the system's start-up phase. During this process, the IP and network address, netmask, and gateway specifications entered by YaST in the <code>/etc/rc.config</code> file are evaluated.
<code>/etc/init.d/route</code>	Sets up static routes over the network.
<code>/etc/init.d/inetd</code>	Starts <code>inetd</code> if specified in <code>/etc/rc.config</code> . This is only necessary for logging in to this machine over the network.
<code>/etc/init.d/portmap</code>	Starts the portmapper needed for the RPC server, such as an NFS server.

Table 2.4: continued overleaf...

<code>/etc/init.d/nfsserver</code>	Starts the NFS server.
<code>/etc/init.d/sendmail</code>	Controls the sendmail process depending on the configuration in <code>/etc/rc.config</code> .
<code>/etc/init.d/ypserv</code>	Starts the NIS server depending on the configuration in <code>/etc/rc.config</code> .
<code>/etc/init.d/ypbind</code>	Starts the NIS client depending on the configuration in <code>/etc/rc.config</code> .

Table 2.4: Some Start-Up Scripts for Network Programs

The User `fwadmin` for the FAS

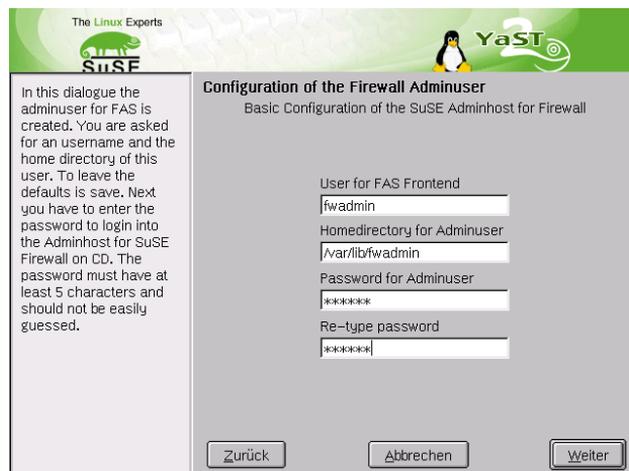


Figure 2.6: The User for the Firewall Administration System (FAS)

In the ‘`fwcdadmin`’ configuration mask (see Figure 2.6), set up a user with which to configure the firewall. Enter a user name. The default name is `fwadmin`. Also specify a home directory for the user. The default for this is `/home/fwadmin`.

Assign a password for the firewall admin user here. This password must be at least five characters in length. In the next screen, define a “pass phrase” and repeat this in the following field (see Figure 2.7).

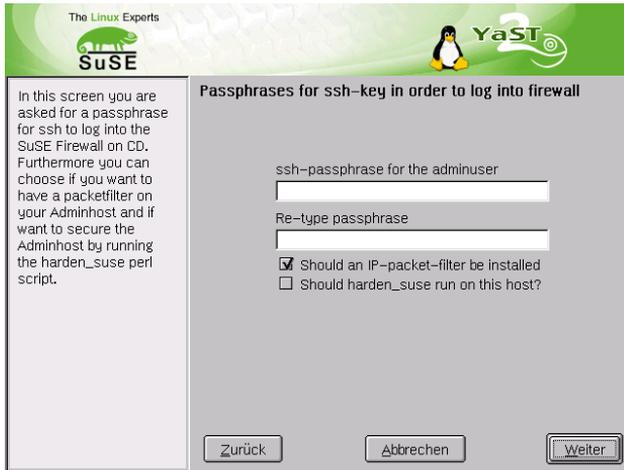


Figure 2.7: Entering an SSH Passphrase for the Admin User

This prepares the basic installation of the SuSE Adminhost for Firewall. Follow the progress of the installation on the screen.

Tip

After the installation, check our Maintenance Web for patches and security updates. Download and install them before configuring the SuSE Firewall on CD. For more information about maintenance and support, refer to [Support, Maintenance, and Patch Management](#) on page 149.

Tip

Upgrade to the VPN Edition

If you purchased the VPN module CD, additional packages that provide VPN functionality for your SuSE Firewall on CD can be installed now from the CD. Log in to the Adminhost and insert the VPN CD. Stop the FAS daemon by starting a shell and, as user `root`, entering the command `rcfasd stop`.

Then start the YcST2 Control Center and select 'Install Patch CD'. Follow the instructions there.

When the installation is finished, restart the FAS daemon with the command `rcfasd start`. The extra VPN functionality is now available.

Firewall Administration System (FAS)

FAS is the graphical administration interface used to create the configuration floppy for the SuSE Firewall on CD. FAS supports multiple users and is able to administer several different configurations. In addition, statistics can be generated and log files evaluated conveniently using FAS.

Logging in as fwadmin	34
Starting the Firewall Administration System	34
Using the FAS	34
Log File Analysis	99
Certificate Management	105
Saving the Configuration	109
Editing an Existing Firewall Configuration	109
Editing Configuration Files	109
Testing the Configuration	110
Documenting Configuration, Tests, and Results	111
Monitoring the Firewall	111

Logging in as `fwadmin`

After installation, the system boots to the graphical login. Log in here as the user `fwadmin` and use the corresponding password. The desktop of the user `fwadmin` opens.

Starting the Firewall Administration System

This is a client and server system, consisting of the GUI and the `fasd` server daemon. The `fasd` (`fas daemon`) manages the various configurations, makes modifications, and checks entries for correctness. The front-end accepts user data and forwards it to the server. Communication between the front-end and the back-end functions in this version via the normal network socket.

Start the FAS from the icon on the desktop (Figure 3.1), by selecting it from the 'K' menu, or by entering `fas` or `FAS` at a command line.



Figure 3.1: FAS Icon

Using the FAS

Initial Login

When logging in for the first time, a dialog appears in which to create a new configuration. Enter the name of the configuration and a comment (see Figure 3.2 on the next page). Then see an overview of the modules.

Creating User Accounts

Logged in as the user `fwadmin`, you can now set up other user accounts. To do this, choose 'Session' → 'Account' → 'Create Login'. A dialog as in Figure 3.3 on page 36) opens.

Enter a new user name and a password. The user name must be at least five characters in length. The password must be between five and eight characters in length. This password protects the configuration of the SuSE Firewall

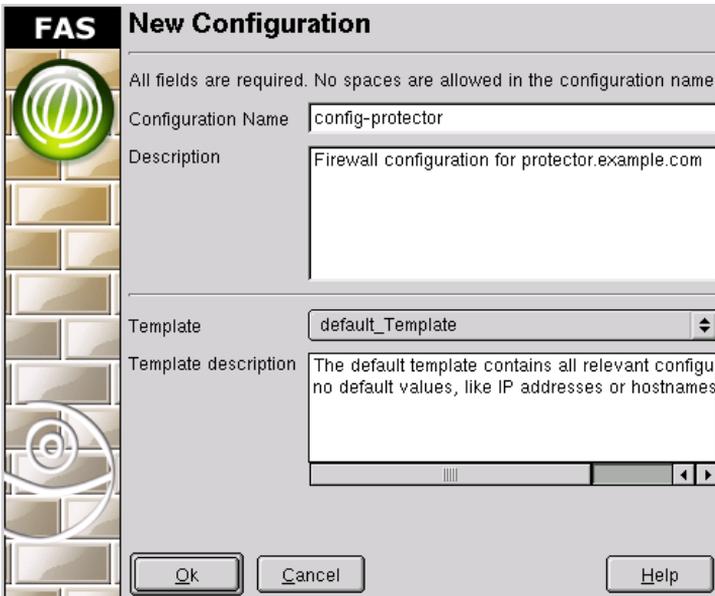


Figure 3.2: Creating a New Configuration

on CD and, for this reason, is checked for its suitability with the program `cracklib`. Select a password that cannot be guessed, so do not use your own date of birth, street name, or the name of your favorite star. It should not be too short, but still be quick to type, so no one can see what you enter. Use a mixture of uppercase and lowercase letters as well as digits.

It is also important to have a password that you can easily remember and do not need to write down. A good method of establishing a password is to make abbreviations of sentences or expressions. Here are some examples:

- NE14TenS (anyone for tennis?)
- AuaEGC (all UNIX admins eat green cheese)
- O10imBd (On the 10th it's my birthday)
- IwihagP (I wish I had a good password)



Figure 3.3: Creating a New Account

Creating a New Configuration

A configuration is created on the initial login. To create additional new configurations, select 'Configuration' → 'New Configuration'. A window appears in which to specify a name for the configuration. Also provide a description of the configuration. This description can be used to explain the purpose of the configuration and to document who has changed what in the configuration when. This configuration description can be extended or edited at any time. Use this option, as good documentation is important. Confirm with 'Ok'.

The new configuration is listed in the left panel. Click the name to display the configuration description in the right panel. Double-click the name to start setting it up.

The left side, as shown in Figure 3.4 on the next page, lists configuration modules — the services with which to configure your firewall — and the status of each module. A red pencil means "must be configured," a red dash means "not active," and a green check mark means "activated." 'Base Setup' and the 'Syslog Module' must be configured first. It is not possible to select any other modules while these modules are in the "must be configured" state. Select a module by clicking it.

The following modules are available:

- Base Setup (network interfaces, routing, host names, root password)
- IP Filter and NAT (configuration of the IP filter rules and network address translation)

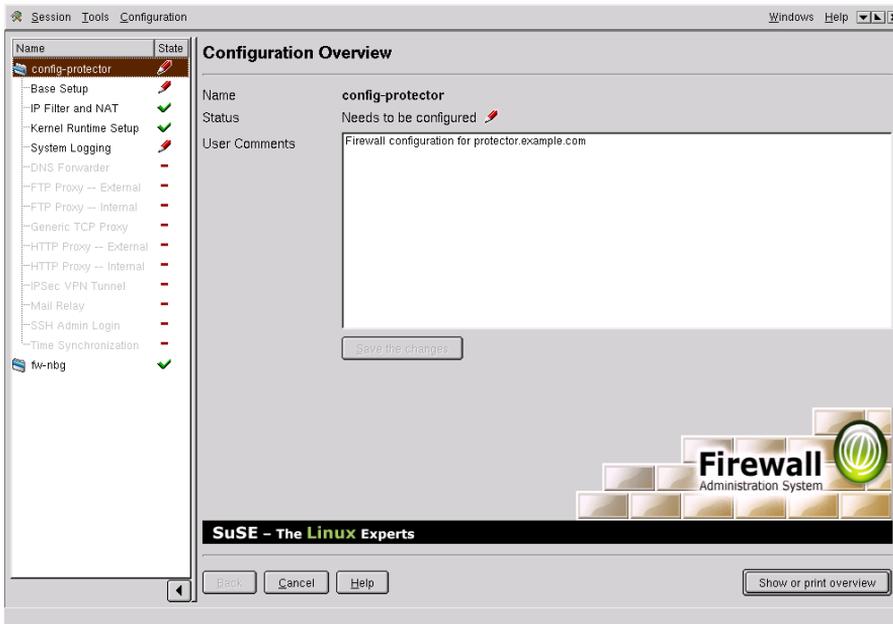


Figure 3.4: Starting a New Configuration

- Kernel Runtime Setup (kernel configuration)
- System Logging (settings for logging)
- DNS Forwarder (name service forwarding)
- FTP Proxy — External (configuration of the FTP proxies from the outside to the inside or to the DMZ)
- FTP Proxy — Internal (configuration of the internal FTP proxies)
- Generic TCP Proxy (configuration of rinetd)
- HTTP Proxy — External
- HTTP Proxy — Internal
- IPsec VPN tunnel (configuration of the additional module for VPN, not included in the Standard Edition)
- Mail Relay (configuration of the mail system)

- SSH Admin Login (login for the administrator with SSH)
- Time Synchronization (configuration of xntpd to synchronize computer time with a time server)

Example, Inc.

This text uses an example configuration of a fictitious company. Only the configuration for the company's headquarters in Nuremberg is described. If the configuration in a branch is significantly different from this, it is pointed out in a short note.

Note

Example firewall configuration

The example firewall used here does not claim to be complete or absolutely correct. It is very important to customize your firewall to your particular situation. Under no circumstances, adopt this example verbatim.

Note

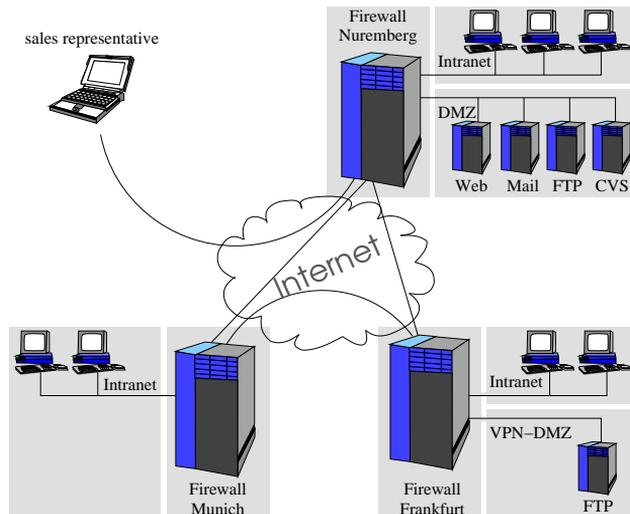


Figure 3.5: Overview of the Example, Inc., Network

The Setup

Example, Inc., an online bookshop with its headquarters in Nuremberg, has 150 staff. It operates branches in Munich and Frankfurt with office and warehouse workers and employs ten sales representatives.

The following infrastructure is required for the business to operate:

- 1 FTP server in the DMZ in Nuremberg
- 1 FTP mirror in the DMZ in Frankfurt
- 1 central web server in the DMZ in Nuremberg
- 1 central mail server in the DMZ in Nuremberg
- 1 CVS server in the DMZ in Nuremberg
- 20 Linux or Windows workstations in Munich
- 50 Linux or Windows workstations in Frankfurt
- 80 Linux or Windows workstations in Nuremberg

The Headquarters in Nuremberg

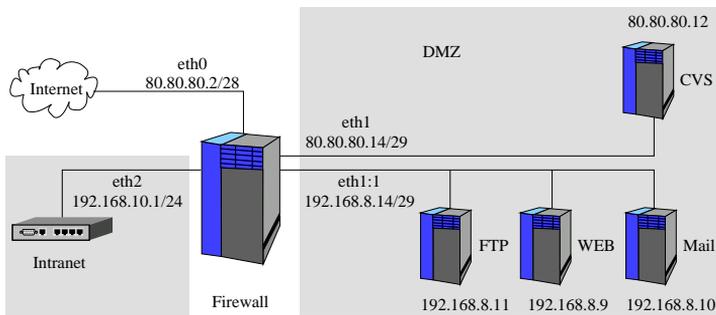


Figure 3.6: The Setup in Nuremberg

In Nuremberg, the bookshop has the entire $80.80.80.0/255.255.255.240$ network. Of this, the network $80.80.80.0/255.255.255.252$ is used as a transfer network between the firewall and the provider. The rest is available for computers in the DMZ. The router of the provider has the IP address

80.80.80.1. The DNS service is available from the provider under the IP addresses 123.123.123.123 and 123.123.123.124. The following entries have already been made:

```
www.example.com = 80.80.80.2 + example.com
ftp.example.com  = 80.80.80.2
mail.example.com = 80.80.80.2 + MX record
cvs.example.com  = 80.80.80.12
```

Because hosts addressed via proxy services of the firewall do not require public IP addresses, two networks are administrated in the DMZ:

```
80.80.80.8/255.255.255.248  services without proxy (CVS)
192.168.8.8/255.255.255.248 services with proxy (web, FTP, mail)
```

The servers in the proxy network in the DMZ receive the IP addresses:

```
web server = 192.168.8.9
mail server = 192.168.8.10
FTP server  = 192.168.8.11
```

The network 192.168.10.0/255.255.255.0 is intended for the connection to the internal network for staff at the headquarters.

The Branch in Frankfurt

In Frankfurt, only a small network is available with real IP addresses. The address space 100.100.100.0/255.255.255.252 also includes the router of the ISP with the IP address 100.100.100.1. The DNS is also managed here by the provider. The following entry was agreed:

```
ftp2.example.com = 100.100.100.2
```

Internally, the network 192.168.11.0/255.255.255.0 is intended for staff.

The Branch in Munich

Munich, which is a small branch, only has a DSL connection to the Internet. There is only one fixed IP address here: 120.120.120.1

The network 192.168.12.0/255.255.255.0 is used internally for the Munich location.

Network Policies

Heads of department in each branch should have full access to the Internet, but all other staff may only have access via proxy services of the firewall. For this reason, the internal networks are again divided. The following setup is agreed for Nuremberg:

A virtual network for heads of department:

```
192.168.10.0/255.255.255.192
```

A virtual network for staff:

```
192.168.10.64/255.255.255.192
192.168.10.128/255.255.255.192
192.168.10.192/255.255.255.192
```

The networks in Frankfurt and Munich are divided accordingly. This results in the following:

- All hosts with IP addresses 192.168.x.1 to 192.168.x.63 have full Internet access.
- Hosts with IP addresses 192.168.x.64 to 192.168.x.254 have Internet access only via proxy services.

In each branch, a DNS server is set up internally with the IP address 192.168.x.65 to answer all internal DNS requests. The DNS servers of the various branches are connected to each other via forward.

The internal domains for the respective branches are:

```
Nuremberg:  nbg-example.com
Frankfurt:  fam-example.com
Munich:     muc-example.com
```

The Adminhost in each branch is set up in the internal network under the IP address 192.168.x.254.

The 10 sales representatives have a basic connection to the headquarters in Nuremberg via the Internet and VPN tunnel. For this, they use winXP or windowsxxx with SSH-sentinel.

Configuring the Base Setup

Click 'Base Setup'. The basic configuration of the firewall is done in five steps:

Basics

In 'Basics', either disable the root password completely (default) or set it. As a safety precaution, repeat the root password (see Figure 3.7). If you do not set a root password, logins directly to the firewall host as root are not possible. Access is then only possible via ssh and RSA keys, assuming SSH is configured.

Check 'Enable serial console' to enable connection of a serial console from which the firewall can be controlled. Also choose between an American or

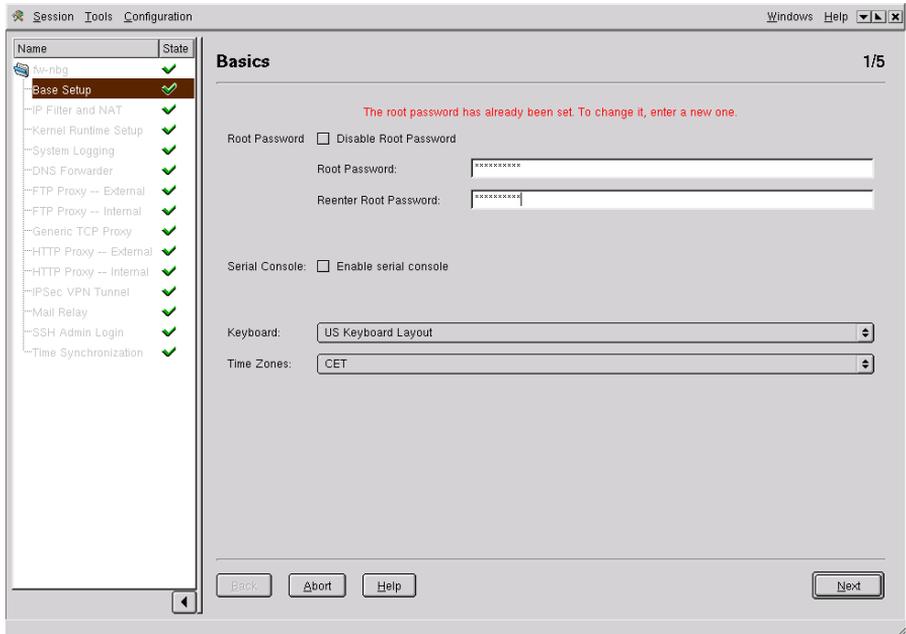


Figure 3.7: Setting the root Password

German keyboard layout. Set the appropriate time zone for the firewall host.

Configuring the Hard Disk

If you activate 'Use hard disk', make the following settings, as shown in Figure 3.8 on the next page.

Disk Device: Hard disk to mount, for example, `/dev/hda` (first hard disk on the first IDE controller)

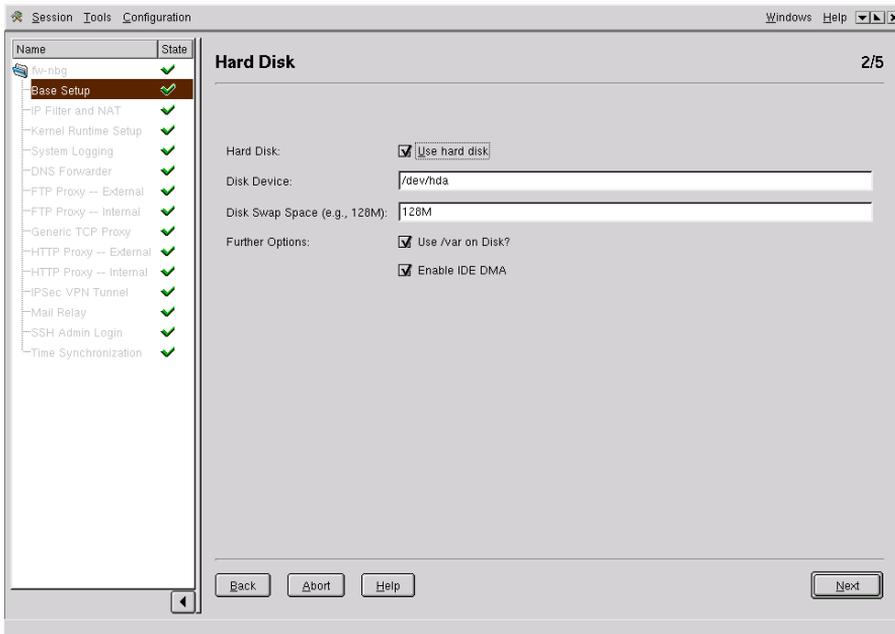


Figure 3.8: Configuring the Hard Disk

Disk Swap Space: Size of the swap partition, such as 128 M

Further Options Activate or deactivate 'Use /var on Disk?' to set whether /var should be on the hard disk.

Enable IDE DMA: Enables DMA for IDE.

If you are using the e-mail proxy, using caching for the HTTP proxy, or want to save messages to the hard disk, the hard disk must be configured and /var activated.

Network Interfaces

If you have already configured network interfaces, find a list of these in this dialog, shown in Figure 3.9 on the following page. At least one internal and one external interface must be specified. Use 'Add' to configure a new interface. In the dialog shown in Figure 3.10 on page 45, select the network type: ethernet (default), ADSL, T-DSL, or ISDN.

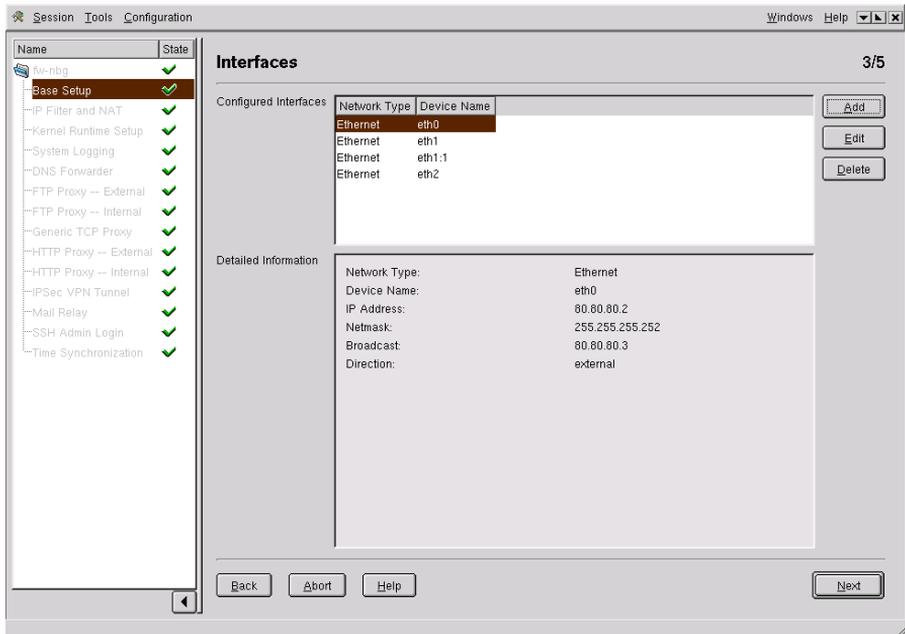


Figure 3.9: Network Interfaces

1. Make these settings in the ethernet dialog (Figure 3.11 on the facing page):

Interface Names are automatically allocated in sequence, but can be changed.

Virtual If required, enter a virtual device name in the field next to this.

IP address IP address to assign to the interface.

Netmask The netmask for the IP address.

Direction internal or external. Is the interface connected to the intranet, to the DMZ, or to the Internet?

Confirm the settings with 'Ok'. Abort with 'Cancel'. The newly configured interface now appears in the list. With 'Edit', modify the configuration of an existing interface. To do this, choose an interface from the list and click 'Edit'. 'Delete' removes the interface selected.

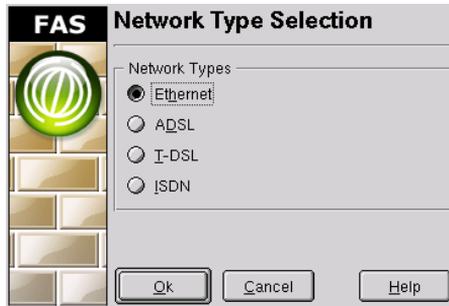


Figure 3.10: Selecting Network Interfaces

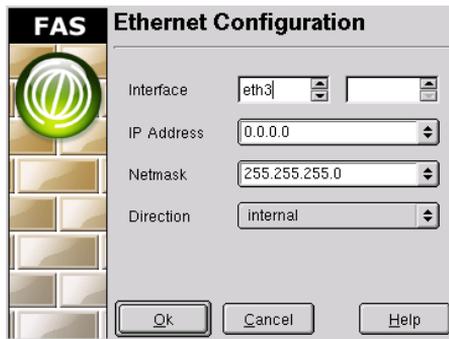


Figure 3.11: Ethernet Interface

- The configuration dialog for DSL is divided into two parts: the DSL configuration and the ethernet configuration. Make the following settings for DSL:

User the name of the user

Password password of the user

Flat rate if you have a flat rate, activate the 'Flat Rate Settings' check box

Idle Time (sec) If you do not have a flat rate, specify the time in seconds after which the DSL connection should be closed if no data is transmitted.

Modify resolv.conf If you activate this check box, the script `modify_resolvconf` temporarily modifies the configuration file `resolv.`

conf for DSL. This file is involved with the resolution of host names by the *resolver* library and contains the domain of the host and the IP address of the name server.

Interface Name Enter the interface for DSL here. ppp0 is the default, but can be changed.

Enter the following settings for the ethernet configuration:

Device The name of the ethernet device

Local IP A local IP address is given here, which can be retained. You only need to change this value if a real network exists with this address.

Remote IP The remote IP address. The default value can also be kept here unless you want to specify a real address. Then you must ensure there can be no collisions.

Close the configuration with 'Ok'.

3. Make the following settings to configure an ISDN card in a dialog like Figure 3.12 on the next page):

ISDN Card Select the appropriate ISDN card from the drop-down menu. Only those cards listed are reliably supported.

Protocol Select the required protocol from the drop-down menu.

Country Prefix The default character + is automatically replaced by the required leading zeros.

Country Code Enter the appropriate country code. 49 is the code for Germany.

Area Prefix Enter a number, if required, to dial before making a local call.

Area Code Enter the area code.

Clicking 'Ok' continues to the second mask of the ISDN configuration. Under 'Dialing Configuration', make the following settings (see Figure 3.13 on page 48):

User Name of the user

Password Password of the user

Idle Time (sec) Enter the time after which the ISDN connection should be stopped if no data is transferred.

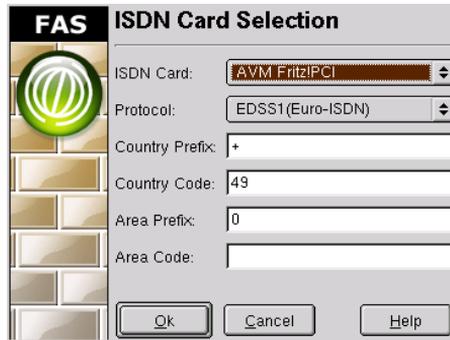


Figure 3.12: ISDN Configuration — Part 1

Interface's Phone Number (MSN) Enter the phone number of the ISDN device (MSN).

Provider's Phone Number Phone number of the provider

Phone Number for dial-in Dial in phone number

Maximum Dial Attempts Maximum number of dial attempts

These details are required under 'Interface Configuration':

Encapsulation Choose here between `syncppp` and `rawip`.

Interface Enter the interface name. The default is `ipp0`, but this value can be changed.

Local IP A local IP address is given here, which can be retained. It is only necessary to change this value if a real network exists with this address.

Remote IP The remote IP address. The default value can also be kept here unless you want to specify a real address. Then you must ensure that there can be no collisions.

Direction Choose the direction: external or internal

Enable multilink Activate this check box if you want to specify a slave interface.

Slave Interface name Enter the interface name for the slave interface.

Click 'Select ISDN Card' if you want to return to the configuration dialog for the ISDN card. To close the configuration, click 'Ok'.

When you have configured all interfaces, click 'Next'.

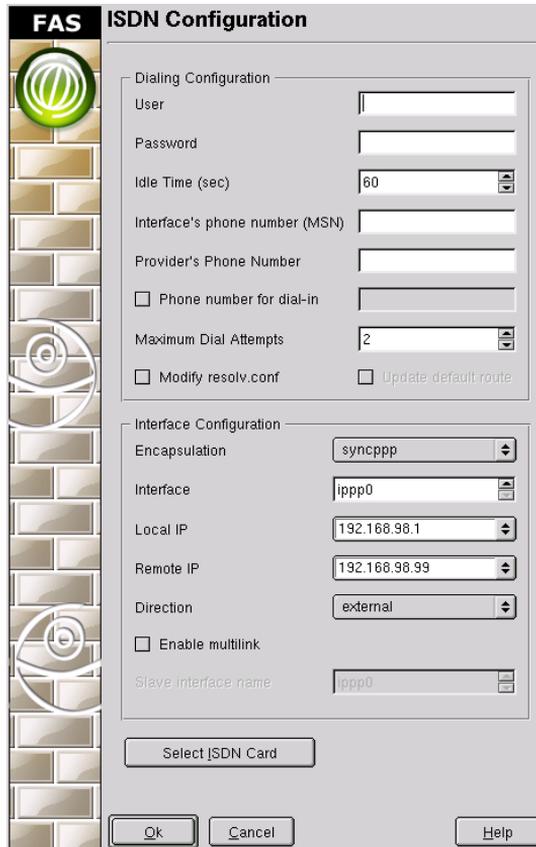


Figure 3.13: ISDN Configuration — Part 2

Routing

In a dialog like Figure 3.14 on the next page), view, create, and modify routes. 'Add' creates a new route. Use 'Edit' to modify an existing route. With 'Delete', remove a set route from the list.

Routes are configured in Figure 3.15 on page 50). The following settings can be made:

Destination Which network or host should the route address? Enter the network address (e.g., 192.168.0.0).

Gateway If desired, enter the IP address for the gateway.

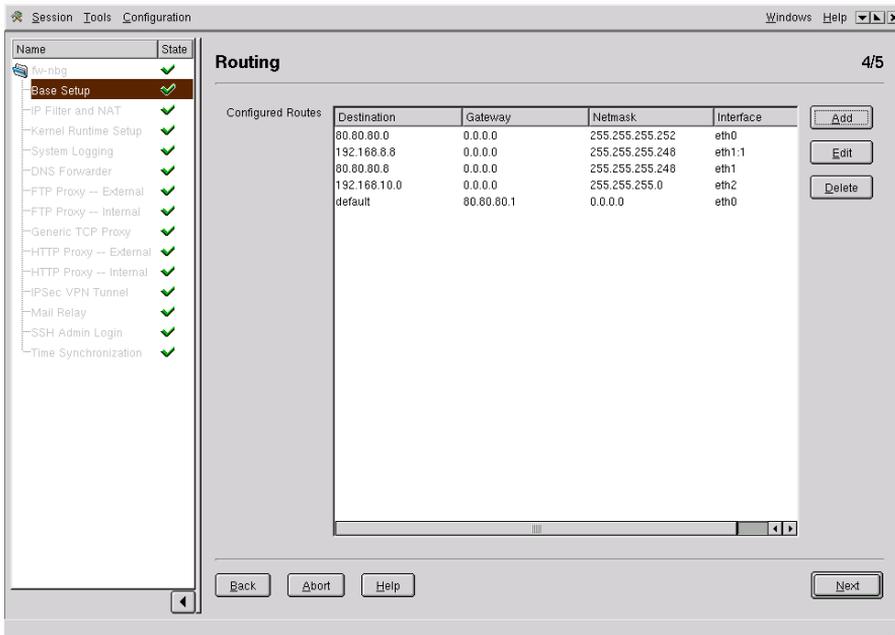


Figure 3.14: Routing Dialog

Netmask The relevant netmask.

Interface Select the interface to use.

Save your settings with 'Ok'. Stop the configuration with 'Cancel'.

Host Name and Name Server Settings

Configure the host name, domain, and name server as shown in Figure 3.16 on page 51.

Firewall Host Name Enter the name of the firewall host. Avoid names like gateway or firewall.

Firewall Domain Enter the name of the domain to which the firewall belongs.

Name server To configure the resolver correctly, the name server and its IP address must be given here. If you press or click 'Add' after entering the IP address, specify additional addresses in the field.

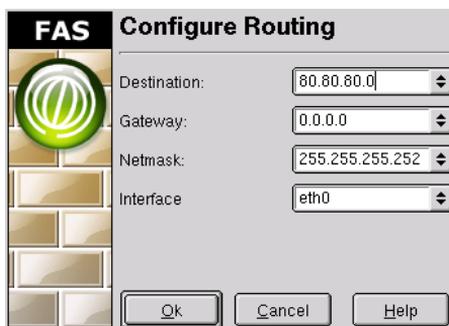


Figure 3.15: Routing

In the next entry line, the search lists are specified, for example, `your-company-inc.com`.

If you now click 'Finish', the base configuration is completed. The message appears briefly on the screen to this effect. If conflicts or problems in the configuration are detected, you will see a corresponding error message here.

The Example, Inc., Configuration

Page 1/5 of the Base Setup

The administrator should be able to log indirectly to the firewall in Nuremberg via console. For this reason the following configuration is set:

```
root password          not selected
```

```
root password:        PvdFSiN2
```

```
confirm root password: PvdFSiN2
```

The keyboard on the firewall has a German layout (qwertz) for our host. This is selected in the drop-down menu for the keyboard.

As the time zone, 'Europe/Berlin' has been chosen. A serial console is not used.

Page 2/5 of the Base Setup

To operate a mail server in Nuremberg, a hard disk must be configured in this menu. First, 'Use Hard Disk' is activated.

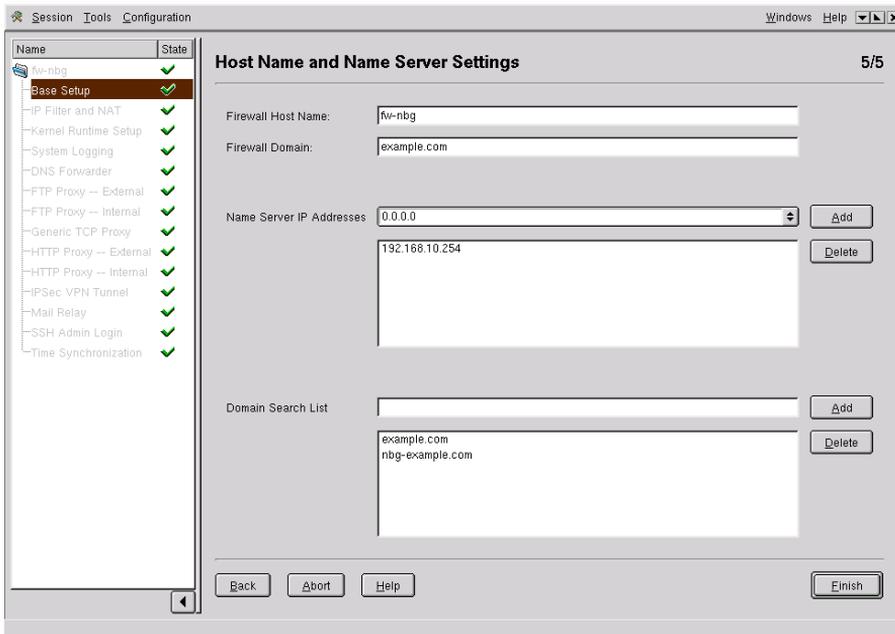


Figure 3.16: Host and Domain Configuration

Only entire hard disks can be used. Individual partitions cannot be configured. In our example, the hard disk is the master on the second IDE bus. Therefore, `hdc` is entered as the device.

A swap space serves to store data from the main memory temporarily on the hard disk, if this is required. 128M is used in our example.

Since an IDE hard disk is being used, it makes sense to activate DMA. This must be supported both by the hard disk and by the hard disk controller. This is irrelevant for SCSI hard disks or RAID arrays.

Page 3/5 of the Base Setup

First, the network card in the direction of the Internet is configured. It is given the device name `eth0`:

```

Network type = ethernet
Device name  = eth0
IP address   = 80.80.80.2
Netmask      = 255.255.255.252
Direction   = external

```

Now the network card (eth1) leading to the DMZ is configured. It should be responsible for a subnet of the public IP addresses. How this subnet is made available externally (announced) is discussed in the Kernel Runtime Setup module.

```
Network type = ethernet
Device name  = eth1
IP address   = 80.80.80.14
Netmask      = 255.255.255.248
Direction    = internal
```

To be able to administer the private subnet in the DMZ, a virtual interface on the same card is also set up:

```
Network type = ethernet
Device name  = eth1:1
IP address   = 192.168.8.14
Netmask      = 255.255.255.248
Direction    = internal
```

Now only the third network card, which should look after the internal network, is missing. It is given the IP address 192.168.10.1:

```
Network type = ethernet
Device name  = eth2
IP address   = 192.168.10.1
Netmask      = 255.255.255.0
Direction    = internal
```

Page 4/5 of the Base Setup

The network routes for the individually configured networks are set by the system itself. At this point, only routes that are not detected by the firewall itself need to be added. In the case of Example, Inc., this is only a default gateway to the router of the provider. The default gateway specifies where packets are sent that do not match the address range of a configured network.

```
Destination = default
Gateway      = 80.80.80.1
Netmask      = 0.0.0.0
Interface    = eth0
```

Page 5/5 of the Base Setup

To complete the base configuration, the host name and domain name for the firewall must be defined.

Firewall host name: `fw-nbg`
Firewall domain: `example.com`

As the name server, the firewall should use the internal DNS server. Then the firewall will also know the names of the internal hosts.

Name Server IP Addresses: `= 192.168.10.65`

In the domain research list, both the publicly known domain as well as the internal domain must be given.

Domain Search List: `example.com`
`nbg-example.com`

IP Filter and NAT

When configuring the IP Filter and NAT module, choose between expert configuration and normal configuration. First, the standard configuration is described.

Click 'Configure'. The IP filter dialog is divided into four masks. Browse through them with the tabs 'IP Forward', 'Masquerading', 'Destination NAT', and 'ICMP to Firewall'.

IP Forward

In the first half of this mask, define a forwarding rule. The following details, also shown in Figure 3.17 on the following page, are required:

Protocol Choose among `tcp`, `udp`, and `icmp`.

Source Address The source address of an IP packet.

Destination Address Specify the destination IP address of the connection.

Destination Port For 'From:', specify the destination port. For 'To:', define a range of ports.

Icmp If you have chosen `icmp` as the protocol, enter the required message type here.

If you then click 'Add', the filter rule appears in the overview window. To edit a rule or delete it, click the rule then the corresponding button.

Under 'Logging', the check box 'Log Access Violation' is activated by default.

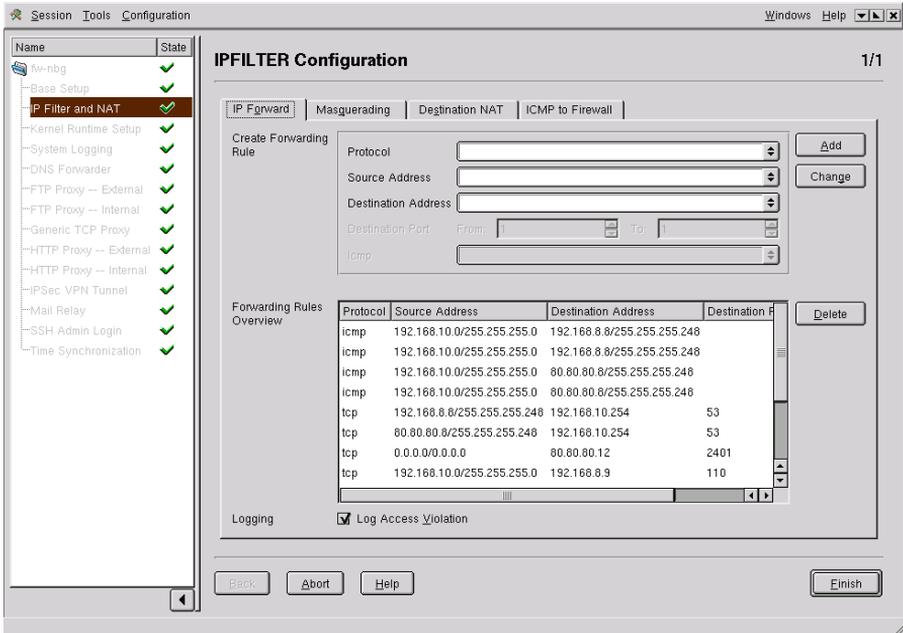


Figure 3.17: IP Forward Dialog

Masquerading

The same entry fields are available in 'Masquerading' (see Figure 3.18 on the next page). Masquerading is a special form of NAT (Network Address Translation). With it, IP packets sent are given the sender address of the router.

Destination NAT

The abbreviation NAT stands for *Network Address Translation*. Destination NAT means that the destination address of the packet is changed. With port forwarding, the destination address of a packet is also changed, for example, if the connection is redirected to a transparent proxy. This dialog is shown in Figure 3.19 on page 56.

Protocol Choose between `tcp` and `udp`.

Source Address Enter the IP address in the local network.

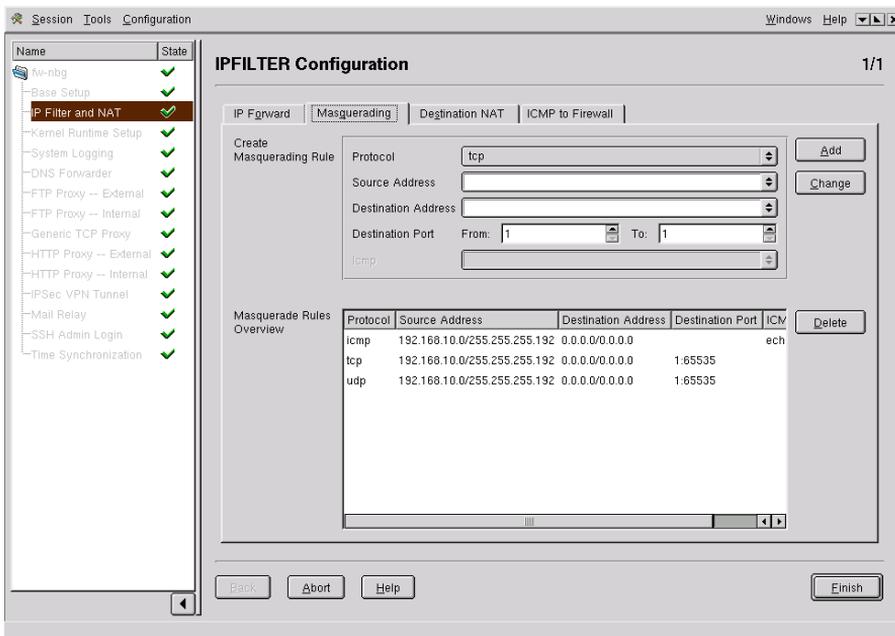


Figure 3.18: Masquerading Dialog

Destination Port For 'From:', enter the destination port. For 'To:', define a series of ports.

Redirect Address Specify the IP address to which the packet is redirected.

Redirect Port Enter the port to which packets are redirected.

ICMP to Firewall

ICMP (Internet Control Message Protocol) is used for error analysis in the network. ICMP send messages describing the error states of IP, TCP, or UDP datagrams. Instead of ports, ICMP has message types containing the header and the first eight bytes of the packet concerned. A well-known example is ping, which sends an echo request to a computer, which then reacts with an echo reply.

In the ICMP mask, shown in Figure 3.20 on page 57, create new rules, with 'Create ICMP Rule', which are then included in the overview window if you click 'Add'. Edit a defined rule by clicking it then editing the entry fields.

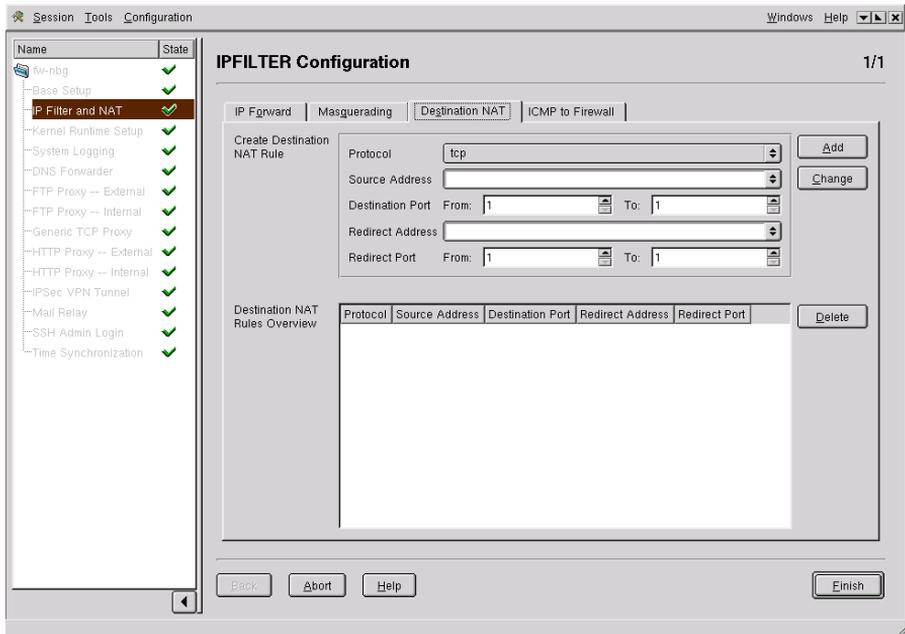


Figure 3.19: Dialog for Destination NAT

Source Address Enter the source IP address.

Destination Address Select the destination address.

ICMP Select the message type.

Logging Activate logging of access violations with the 'Log Access Violation' check box (on by default).

The Example, Inc., Configuration

IPFilter Configuration

This is the most time-consuming part of the configuration. In this module, you should know exactly what effects the entries have.

1. IP Forward: All the rules the FAS cannot generate automatically must be entered here. These are all rules between two networks that are not covered on the firewall by a proxy service.

The following rules are required for Example, Inc.:

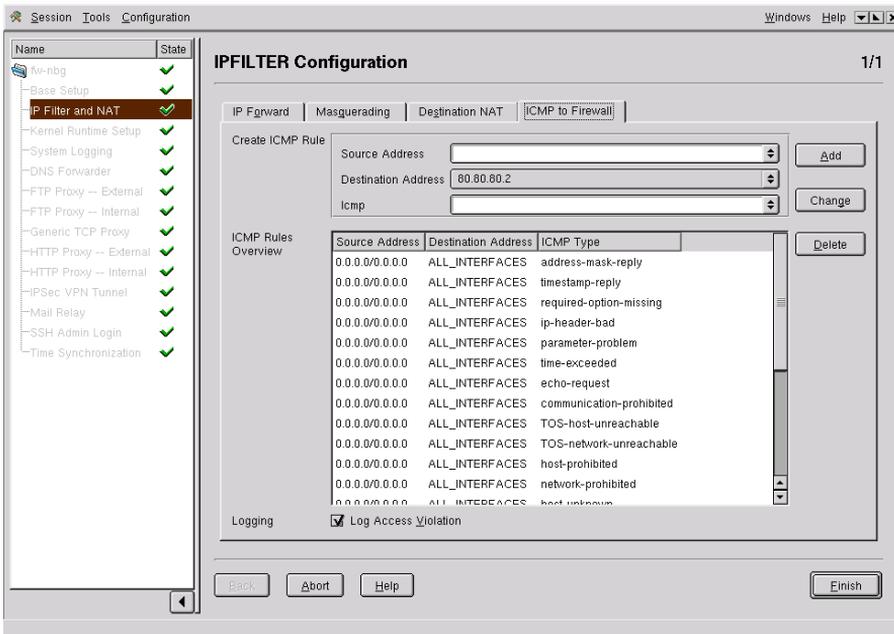


Figure 3.20: Dialog for ICMP

- Address the internal DNS server from the DMZ

Protocol	UDP
Local address	192.168.8.8/255.255.255.248
Remote address	192.168.10.65
from Port	53
to Port	53
Protocol	UDP
Local address	80.80.80.8/255.255.255.248
Remote address	192.168.10.65
from Port	53
to Port	53

- make the CVS server in the DMZ accessible from everywhere

Protocol	TCP
Local address	0.0.0.0/0.0.0.0
Remote address	80.80.80.12
from Port	2401
to Port	2401

- POP3 access to the mail server for clients from the internal network

Protocol	TCP
Local address	192.168.10.0/255.255.255.0
Remote address	192.168.8.10
from Port	110
to Port	110
- IMAP access to the mail server for clients from the internal network

Protocol	TCP
Local address	192.168.10.0/255.255.255.0
Remote address	192.168.8.10
from Port	143
to Port	143
- SMTP access to the mail server for clients from the internal network

Protocol	TCP
Local address	192.168.10.0/255.255.255.0
Remote address	192.168.8.10
from Port	25
to Port	25
- LDAP access to the mail server for clients from the internal network

Protocol	TCP
Local address	192.168.10.0/255.255.255.0
Remote address	192.168.8.10
from Port	389
to Port	389
- ICMP messages from the internal network to the DMZ (not mandatory, but desired at Example, Inc.)

Protocol	ICMP
Local address	192.168.10.0/255.255.255.0
Remote address	192.168.8.8/255.255.255.248
Protocol	ICMP
Local address	192.168.10.0/255.255.255.0
Remote address	80.80.80.8/255.255.255.248

2. Masquerading

All heads of department should be granted full access to the Internet.

Protocol	TCP
Local address	192.168.10.0/255.255.255.192
Remote address	0.0.0.0
from Port	1
to Port	65535
Protocol	UDP
Local address	192.168.10.0/255.255.255.192
Remote address	0.0.0.0
from Port	1
to Port	65535
Protocol	ICMP
Local address	192.168.10.0/255.255.255.192
Remote address	0.0.0.0

3. Destination NAT Destination NAT rules are not required at Example, Inc.

Kernel Runtime Setup

The Kernel Runtime Setup is a matter for professionals. By default, sensible values have been set there for the SuSE Firewall on CD. The module is shown in Figure 3.21 on the following page. In most cases, the various modules access the relevant entries automatically. Documentation about this can be found in the kernel documentation in the kernel source package. Do not change anything here if you are not completely sure of the implications.

The Example, Inc., Configuration

Example, Inc., received a large network from its provider (255.255.255.240). Since no additional transfer network is available between the router and the firewall, the firewall must distribute the public addresses of the DMZ in the direction of the router.

To achieve this, the `proxy-arp` function is activated on the network interface `eth0`. This causes the kernel to reply to all arp requests for routes known to it. To do this, select `'net' → 'ipv4' → 'conf' → 'eth0' → 'proxy_arp'` then activate `'Status'`.

By default, the kernel has a delay in responding to `proxy_arp` requests. This should be reduced with `'net' → 'ipv4' → 'neigh' → 'eth0' → 'proxy_delay' → '10'`

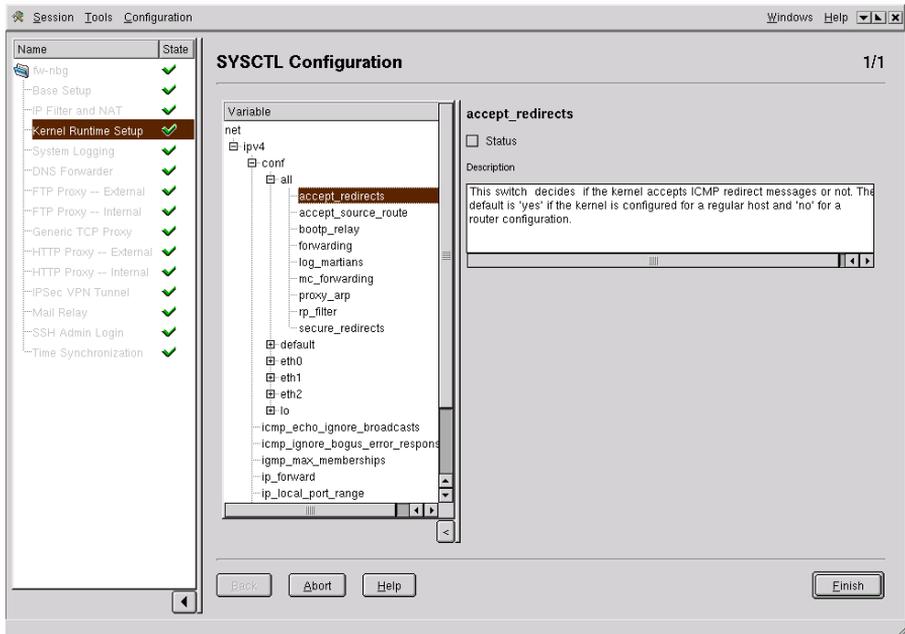


Figure 3.21: Kernel Runtime Settings

System Logging

In this dialog, shown in Figure 3.22 on the next page, configure the behavior of the Syslog daemon. You have the option of sending the output of the syslogd to the text console 'Log to /dev/tty9' (Ctrl + Alt + F9), to the hard disk 'Log to hard disk', and to a log host. Enter one or more IP addresses of hosts to which the logs should be written. These hosts must be configured to accept the logs. The Adminhost is already prepared for these tasks.

Log evaluation in FAS is activated if you enter the IP address of the Adminhost and select 'Enable Log and Traffic Evaluation'. Evaluating logs using FAS is described in [Log File Analysis](#) on page 99.

The Example, Inc., Configuration

Syslog Module

Example, Inc., would like to use all logging possibilities of the firewall. For this reason, all control boxes are activated. Adminhost is given as the desti-

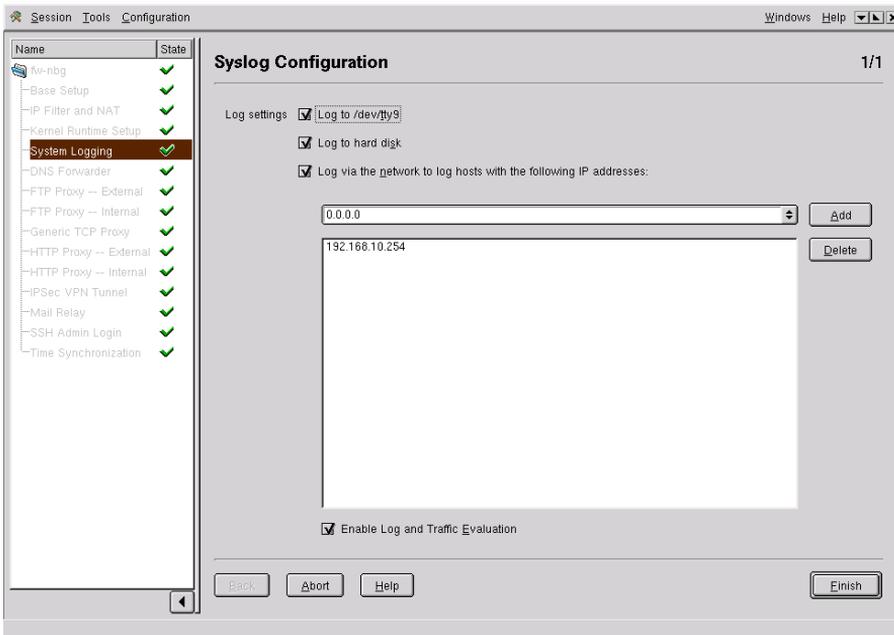


Figure 3.22: Settings for Syslog Daemon

ation for the log files: 192.168.10.254. 'Enable Log and Traffic Evaluation' is activated.

DNS Forwarder

In this dialog, shown in Figure 3.23 on the next page, name service forwarding is configured. Detailed explanations of DNS and `bind8` can be found in [DNS — Domain Name Service](#) on page 157. Enter the IP address of the host to which name service requests should be forwarded. If you activate 'Forward only', requests are sent *exclusively* to the hosts in the list. If this option is not activated, requests are first be sent to the hosts specified then, if necessary, to the root name server on the Internet.

In the lower half of this dialog, assign IP addresses or interfaces for `bind`. `bind` only accepts requests when they arrive on these interfaces. Alternatively, activate 'Auto listen-to:'.

In the second dialog, shown in Figure 3.24 on page 63, select for IP filter

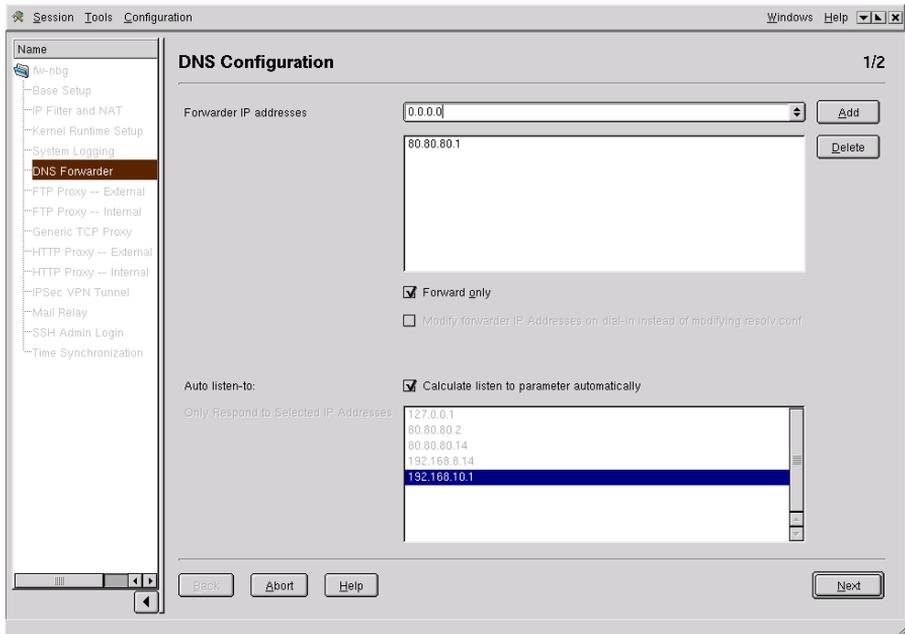


Figure 3.23: DNS Configuration

rules generated automatically and specify the IP addresses of the hosts to write to the file `hosts.allow`.

The Example, Inc., Configuration

Page 1/2 of the DNS Configuration

The proxy for DNS is only set up for an internal host. This should be the internal name server.

Two name server addresses have been made available by the provider, which are used as forwarders. Requests will be forwarded to these name servers.

```
Forwarder: 123.123.123.123
           123.123.123.124
```

As it can be assumed that these names servers will always be active and that no names should be resolved directly, 'Forward only' is activated.

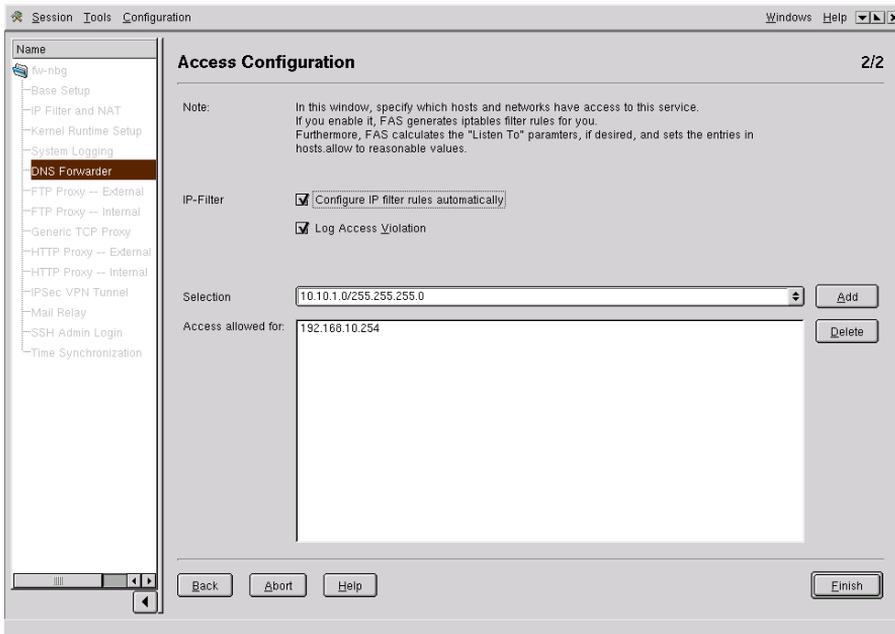


Figure 3.24: DNS Access Configuration

Page 2/2 of the DNS Configuration

'Configure IP filter rules automatically' must be activated for the necessary `iptables` rules to be generated. Now, which IP addresses may use the DNS proxy should be defined. In Example, Inc., in Nuremberg, this is only the internal DNS server.

Access allowed for: 192.168.10.65

FTP Proxy — External

If you are operating your own FTP server, make the corresponding settings here, as in Figure 3.25 on the following page).

Auto listen to Interface on which the FTP server accepts connections.

Choose between the 'Auto listen to' function or one of the possible IP addresses.

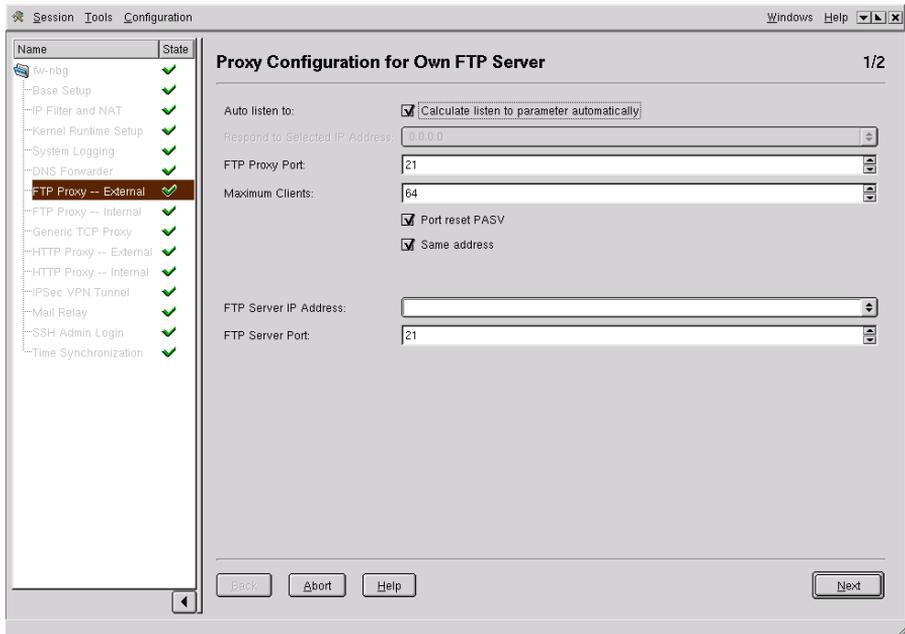


Figure 3.25: Configuration of the FTP Server

FTP proxy port Port on which the FTP proxy is addressed, normally port 21. Viewed from the outside, this turns your firewall into the FTP server. Ask your provider to make an alias entry in the DNS for your firewall, for example, ftp.mycompany.com

Maximum Clients Maximum number of FTP clients that can be connected to the FTP server at any one time.

Port reset PASV By default, this check box is activated so passive FTP mode is selected. In passive mode, the FTP client asks the server which ports to use. The client then opens this port for data transfer. If you deactivate the check box, data transfer takes place in active mode, which means the client sends a request to the FTP server, which then opens a port for data transfer to the client.

Same address The check box is activated by default. In this case, the IP address to which a request is sent must match the address from which a reply comes.

IP address of the FTP server Enter the IP address of the FTP server located in the intranet or in the DMZ.

FTP server port Port on which the FTP server is listening. Normally, this is port 21.

Click 'Next' to continue to the second mask of the dialog, illustrated in Figure 3.26. Have the IP filter rules generated automatically by activating the corresponding checkbox. In the selection field, enter the IP addresses and networks that should have access to the FTP service. With 'Finish', complete the configuration.

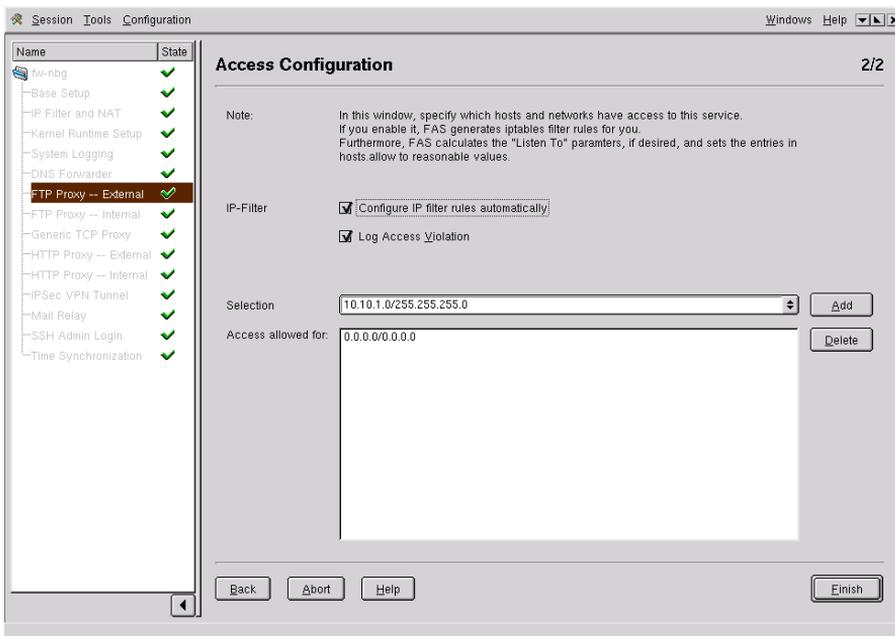


Figure 3.26: Access to the FTP Server

The Example, Inc., Configuration

Page 1/2 Proxy Configuration for the FTP Server

Although the FTP server of Example, Inc., has a public IP address, it should be protected by a proxy service. Here, the filter rules should also be generated automatically. For this reason, 'Auto listen to' is activated. The FTP proxy should run on the standard FTP port 21.

You will not expect many clients to want access to the FTP service at the same time. To avoid an overloaded line, the maximum number of clients who can be logged in at the same time is restricted to 20.

The options 'Port reset PASV' and 'Same address' must only be modified in special cases and should remain activated. Finally, make the address of the real FTP server known to the firewall:

FTP Server IP Address: 192.168.8.11
FTP Server Port: 21

Page 2/2 Proxy Configuration for FTP Server

Setting the IP filter rules should again be done automatically, so 'Configure IP filter rules automatically' is activated. As there are no restrictions in the access to the FTP server, the IP address should be specified as follows:

Access allowed for: 0.0.0.0/0.0.0.0

Caution

Public access

The address 0.0.0.0/0.0.0.0 overrides all access restrictions for the FTP proxy. This means that any machines can send requests to the FTP proxy and server.

Caution

FTP Proxy — Internal

In this dialog, specify which users of the intranet may access FTP servers in the Internet via which connections. The dialog is shown in Figure 3.27 on the facing page.

Auto listen to Decide if the 'listen-to' parameters should be generated automatically (activated by default).

Reply to Selected IP Address If 'Auto listen to' is switched off, enter an address here via which the FTP proxy is addressed.

FTP proxy port 10021
Port on which the FTP proxy is addressed.

Maximum Clients Maximum number of open connections on the FTP proxy.

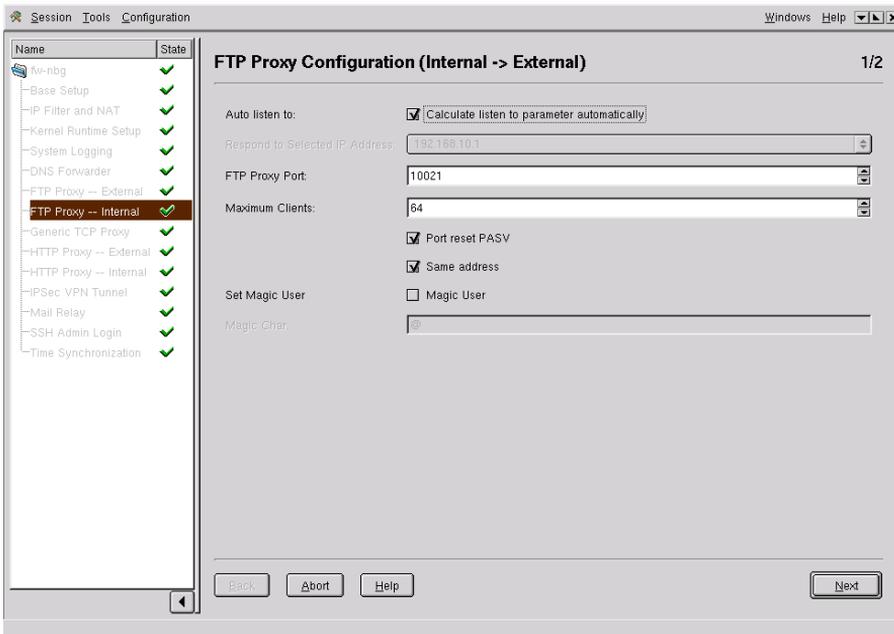


Figure 3.27: Configuration of the Internal FTP Proxy

Port reset PASV By default, this check box is activated, so passive FTP mode is selected. In passive mode, the FTP client asks the server which ports to use. The client then opens this port for data transfer. If you deactivate the check box, data transfer takes place in active mode, which means the client sends a request to the FTP server, which then opens a port for data transfer to the client.

Same address The check box is activated by default. In this case, the IP address to which a request is sent must match the address from which a reply comes.

Magic User: If you activate this check box, the selected destination FTP server is given the user name `user@host:port`. This may then appear as follows:

```
> ftp user@remoteftp.remote.org:21
```

Magic Char: The Magic Char character is set to `'%'` by default. If the option 'Magic User' is activated, any character can be chosen.

In the second mask, shown in Figure 3.28), configure access to the internal FTP proxy. By default, automatic generation of filter rules is activated. Select the IP addresses that should be given access. Select 'Finish' to complete the configuration.

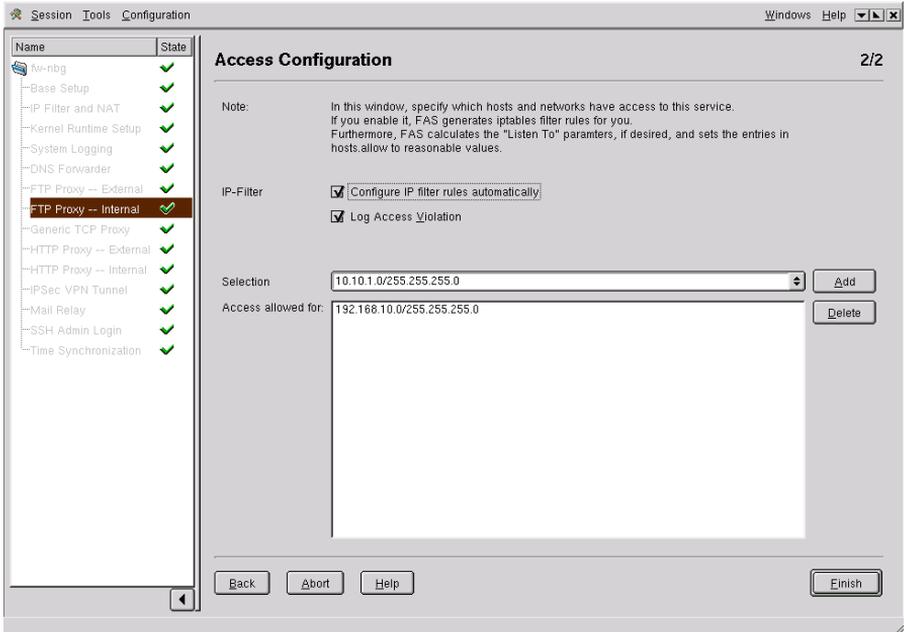


Figure 3.28: Access to the Internal FTP Proxy

The Example, Inc., Configuration

Page 1/2 Internal FTP Proxy Configuration

Only a few basic settings are made here. As usual, the corresponding rules are generated automatically.

Auto Listen to:	activated
FTP Proxy Port:	10021
Maximum Clients:	64
Port reset PASV	activated
Same Address	activated

As the FTP proxy is transparent and for intranet users, no 'Magic User' is required.

Page 2/2 Internal FTP Proxy Configuration

In this module, the filter rules are also generated automatically. The internal network of Example, Inc., is added for the allowed clients:

```
Access allowed for: 192.168.10.0/24
```

Generic TCP Proxy

`rinetd` is used as the generic proxy. This is software that accepts a connection on one interface and forwards the incoming data with another interface to a different machine. This is port-dependent. It is really the routing of TCP connections on the application level.

`rinetd` can only route connections across one channel. It cannot be used as an FTP proxy because an FTP connection uses two channels.

The generic proxy `rinetd` should be used if there is no dedicated Application Level Gateway (such as `ftp-proxy-suite` or `Squid`). With `rinetd`, for example, direct connections for `pop3` through the firewall in a simple and secure manner.

`rinetd` also supports complete logging — all incoming connections are recorded by `syslogd`.

Read more in the man page for `rinetd` (`man rinetd`).

Configuring the Generic Proxy

The ‘Forwarding’ dialog is shown in Figure 3.29 on the following page. Click ‘Add’ to make the settings for `rinetd` connections.

In the first dialog (Figure 3.30 on page 71), enter the following data:

Bind address is the IP address on which `rinetd` should accept a connection.

Bind port is the port number of the service to forward.

Connection address is the IP address of the host providing a service.

Connection port is the corresponding port number for the service.

Click the ‘RINETD ACLs’ tab to define the “allow” and “deny” rules. These rules restrict the source addresses of incoming TCP connections. Specify individual IP addresses or IP address ranges. Allowed characters in the allow and deny attributes are the digits 0 to 9, period (‘.’), question mark (‘?’),

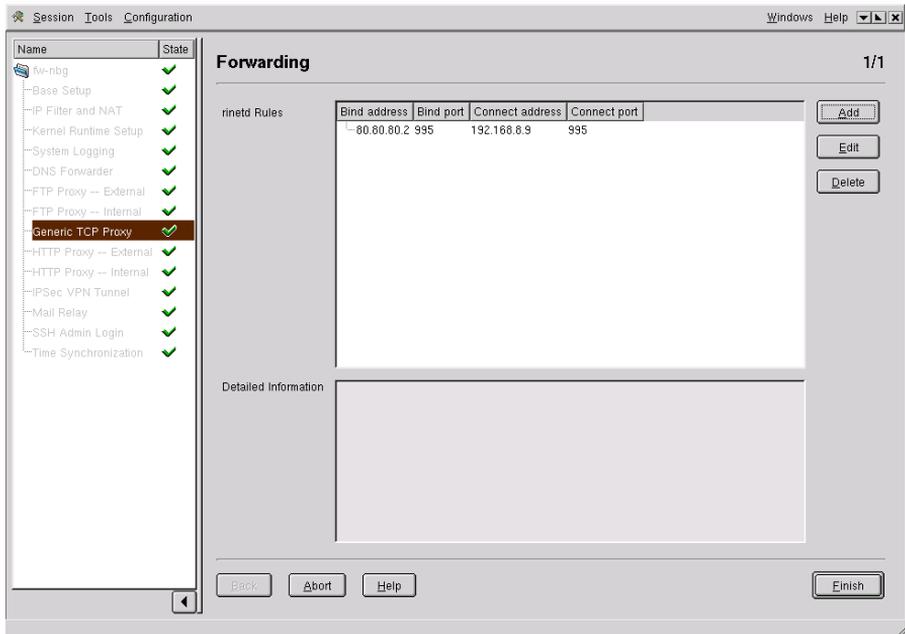


Figure 3.29: Configuration of the Generic Proxy

and asterisk ('*'). The wild card '?' stands for any character at all. Asterisk '*' represents any number of any characters.

If one or more allow attributes are set for a rule, all connections that do not match any of these attributes are immediately rejected. The same thing is the case for the deny attributes. Incoming TCP connections from a source address that matches a deny rule are immediately rejected. Others are allowed.

Click the 'IP Filter Configuration' tab. The automatic IP filter configuration is activated by default. In addition, you can activate logging for "Access Violations" by checking it. Then select those IP addresses that should have access to rinetd and click 'Add'.

To edit an existing rule, select the rule from the list and click 'Edit'. A dialog appears in which the current settings are visible. Make changes and confirm 'Ok'. To delete a rule, select it and press 'Delete'. With 'Finish', complete the configuration of the generic proxy.



Figure 3.30: Redirect Settings for `rinetd`

The Example, Inc., Configuration

Generic TCP Proxy

In Example, Inc., new sales representatives do not always have access to the intranet from the beginning. For them to be able to read and send mail when away, the services `pop3s` and `imap3` are set up on the mail server.

1. Redirect Settings

Bind address	80.80.80.2
Bind port	995
Connect address	192.168.8.10
Connect port	995

2. RINETD ACLs

All sales representatives connect to the provider of Example, Inc. The provider allocates its dynamic addresses in the network `80.80.60.0/255.255.255.0`. Access can therefore be restricted to this network, creating additional security.

allow activated
Pattern 80.80.60.*

3. IP Filter Configuration

No further restrictions need to be made in the “IP Filter Configuration” tab. Only the standard access remains to be configured:

IP Filter	activated	Configure IP Filter rules automatically
	activated	Log access violation
Access allowed for:	80.80.60.0/24	

In addition, a second similar rule must be inserted. The only difference from the previous rule is in the port number. Both `Bind port` and `Connect port` are set to the value 993 for `imaps`.

Configuring the HTTP Proxy — External

If you are not operating your own web server, you do not need to configure anything here. Otherwise, make settings for the HTTP proxy in this module. The module processes HTTP requests from the outside. The following entries, shown in Figure 3.31 on the facing page, are required to configure the HTTP proxy:

HTTP proxy port Port on which the proxy should accept HTTP requests from the outside. The default is port 80.

Auto listen to Decide if the “listen to” parameters should be calculated automatically or if you want to set specific IP addresses on which the HTTP proxy accepts requests.

Listen to selected IP addresses: Specify the IP addresses of the interfaces from which the proxy accepts HTTP requests.

Create Redirect Rule If an HTTP request arrives on the firewall, this is redirected to the corresponding web page. Enter the local path in the first field and, in the second field, the URL for the redirection.

When all settings have been made, confirm with ‘Next’.

In the following dialog (see Figure 3.32 on page 74), define who has access to the HTTP proxy. Automatic generation of IP filter rules is activated by default. If required, select the logging of “Access Violations”. Then choose the IP addresses to have access to the proxy. With ‘Add’, these addresses are applied. With ‘Finish’, complete the configuration.

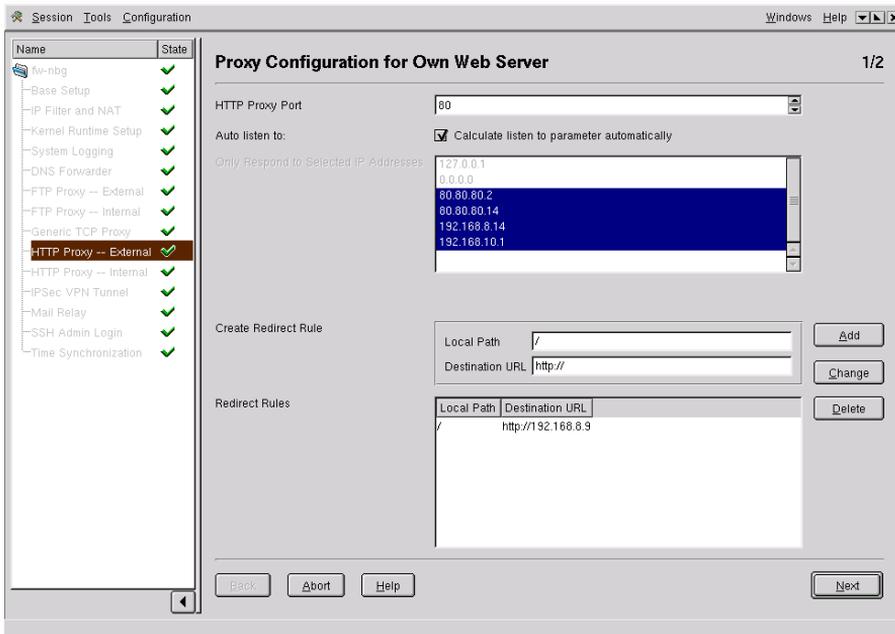


Figure 3.31: Configuration of the External HTTP Proxy

The Example, Inc., Configuration

Page 1/2 HTTP Proxy — External

The external HTTP proxy has to know on which port the Web server can be accessed.

HTTP Proxy Port 80

The basic configuration of the firewall rules can again be generated automatically by activating 'Auto listen to'.

All that is left are the forwarding rules for the proxy. In this way, requests to <http://www.example.com/> are forwarded to the web server. The following rules are necessary for this:

```
local path:  Destination URL:
/           http://192.168.8.9
```

Page 2/2 HTTP Proxy — External

In 'Access Configuration', it is possible to restrict the networks from which requests may be made. This is not wanted here so 'Access allowed for' is set

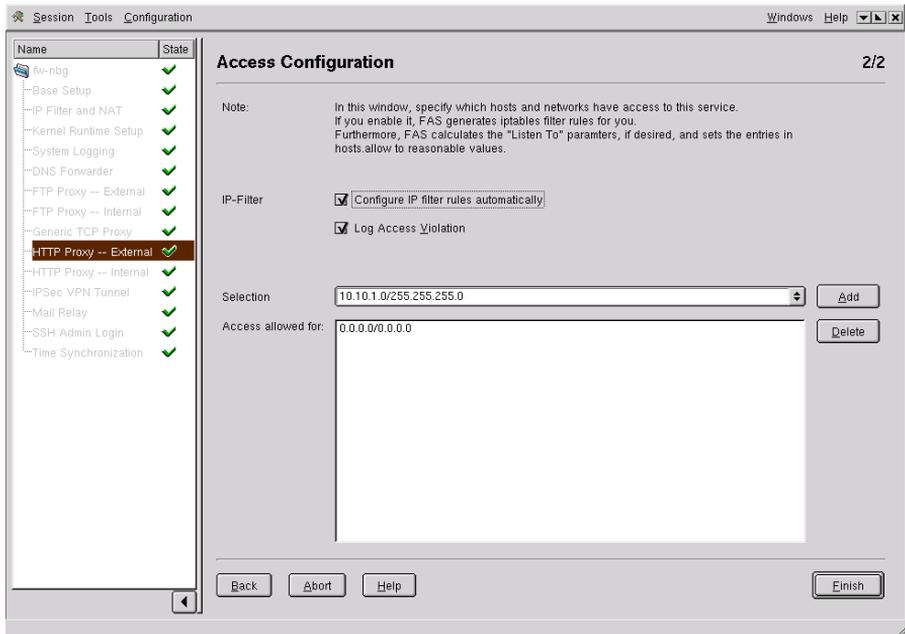


Figure 3.32: Access to the External HTTP Proxy

to 0.0.0.0/0.0.0.0.

Configuring the HTTP Proxy for Internal Connections

In the 'HTTP Proxy — Internal' module, configure the proxy for the users of the internal network. This dialog is shown in Figure 3.33 on the facing page.

HTTP Proxy Port The default HTTP port for internal requests is set to 3128.

Auto listen to Decide if the "listen to" parameters should be automatically calculated or if you want to specify certain IP addresses on which the HTTP proxy accepts requests.

Listen to selected IP addresses Here, specify the IP addresses of interfaces on which the proxy accepts HTTP requests. In the selection list, find the interfaces marked in the base configuration as internal. 0.0.0.0 stands for all interfaces of the firewall.

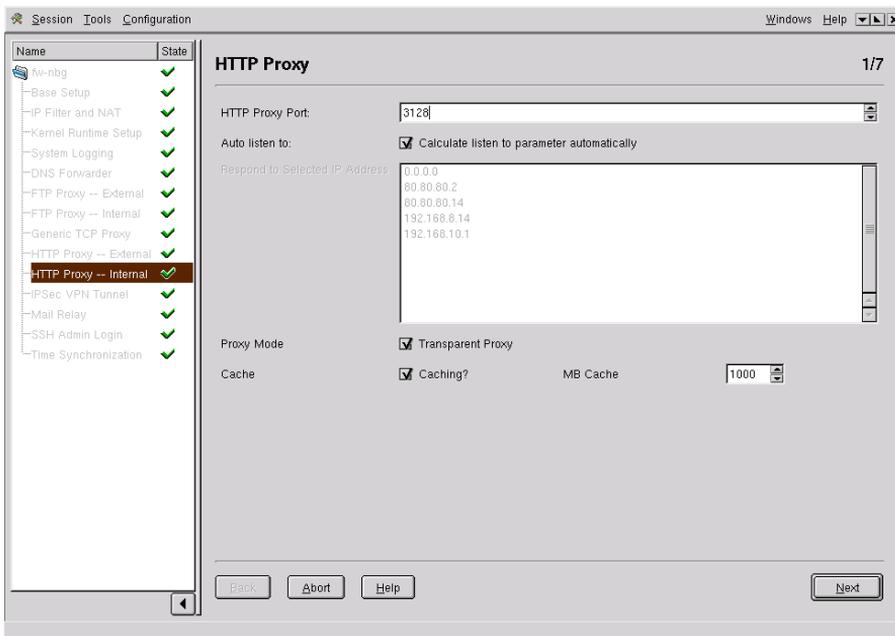


Figure 3.33: Configuration of the Internal HTTP Proxy

Transparent Proxy Because HTTP requests from clients normally arrive on port 80 so are ignored by the proxy, these requests can be redirected to port 3128. For this, ‘Transparent Proxy’ must be activated.

Caching If HTTP requests are repeated, activate the option ‘Caching?’ to avoid duplicate processing of valid pages. Specify the size of the cache in the corresponding entry field. It should not be less than 100 MB.

When you have finished all settings, confirm with ‘Next’.

Defining ACLs

As shown in Figure 3.34 on the next page, define “Access Control Lists”, which grant or refuse specific users access to specific web pages. The rules are processed from top to bottom until there is a match.

Name for ACL First, give a name for the list to create.

Type of ACL Next, choose a ‘Type’ for your ACL. Choose from:

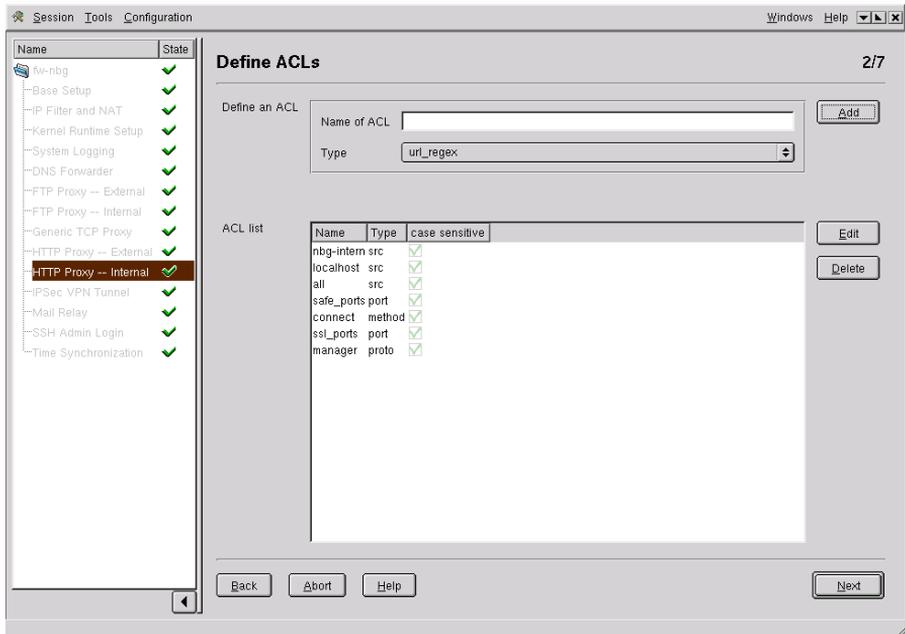


Figure 3.34: Define ACLs

url_regex Details of URL address using regular expressions.

proto (protocol) Specify the appropriate protocol here.

src (source) Defines source addresses.

dst (destination) Details of destination addresses.

port Details of the port.

method Details of the method, such as CONNECT, POST, or GET.

srcdomain Name of the source domain.

dstdomain Name of the destination domain.

srcdom_regex Name of the source domain using regular expressions.

dstdom_regex Name of the target domain using regular expressions.

time Details of time.

urlpath_regex Details of URL paths using regular expressions.

browser Names of browsers.

maxconn Defines the maximum number of connections.

Add Add the new ACL to the list of already created ACLs.

Edit If you click 'Edit', a window opens in which to enter or modify values valid for the ACL selected.

Delete Delete ACLs with this button.

When you have finished the entries, confirm with 'Next'.

Arranging ACLs

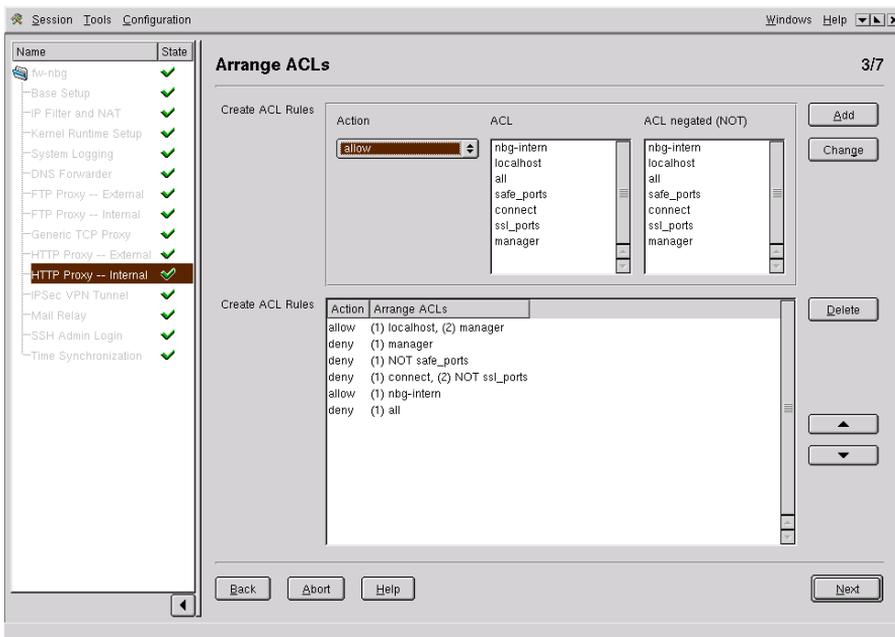


Figure 3.35: Arrange ACLs

Find various ways of setting up rules (ACL) on this module page, shown in Figure 3.35. The settings for 'ACL' and 'negate ACL' are linked with AND.

- With 'Action', choose between 'allow' and 'deny' for Internet accesses, which are defined below.
- Define, via 'ACL', an already created list for which the settings are made valid.

- Define, via 'Negate ACL', an ACL to use in negated form.
- A new rule is integrated with 'Add'. It is then be listed in the list field of defined ACLs.
- With the black arrow keys, move a selected rule up or down in the list.
- To edit rules, click the corresponding rule in the list field. The settings appear in the upper part. Modify these and apply the changes with 'Change'.
- To delete rules, select the appropriate rule and click 'Delete'.

Caution

The order of the rules in the list is very important, because the list is processed from top to bottom. Depending on what matches first, access to the requested URL is either given or denied.

Caution

- When you have made all changes, confirm the settings with 'Next'.

Content Filter

This module is shown in Figure 3.36 on the next page. To filter HTML page contents or to search for specific contents and possibly block them, activate 'Content Filter'.

Note

You should be able to write HTML code to configure this filter. The tags and attributes mentioned below are components of HTML, the description language in which web pages are written.

Note

All filter settings created can be found in the lower part of the window. HTML tags and attributes that are not defined are rejected.

Proceed as follows:

- To create a filter rule, select the appropriate type from the selection field with the same name. You have a choice of 'tag' and 'attr' (attribute). In 'Name', enter the text of the tag or attribute then select an action. It is possible to specify multiple actions. You have the choice of:

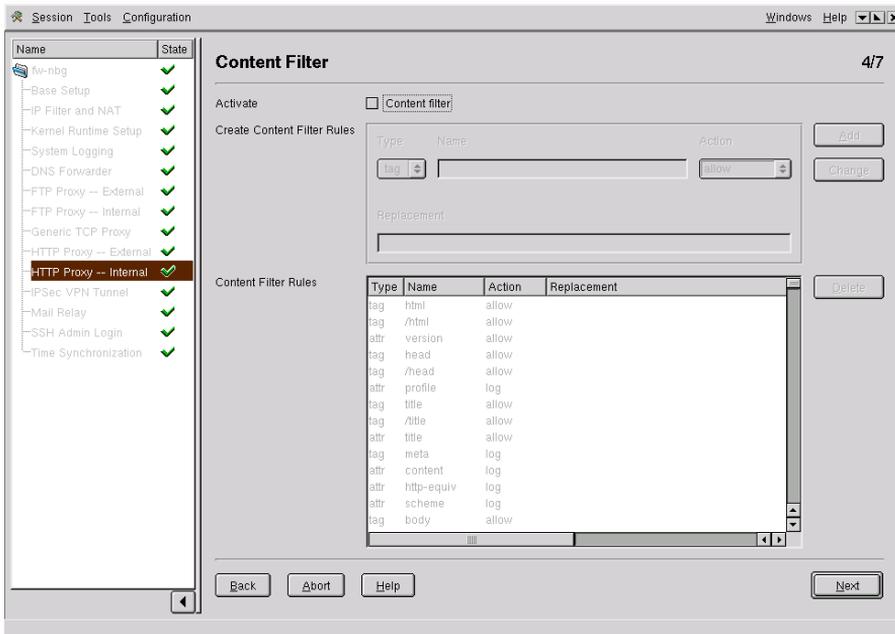


Figure 3.36: Content Filter Dialog

allow allows the selected tag/attribute.

log creates an entry in the log file, which can be later used for creating filter strategies.

log,allow creates an entry in the log file and allows the selected tag attribute.

replcont replaces the tag or attribute and continues to evaluate it.

replabort replaces the tag or attribute then stops evaluating it.

- To remove entire code parts from requested documents, the following two actions are used:

replskip shows the beginning of a deletion that is not forwarded to the user.

replendskip shows the end of the deletion.

- In 'Replace with', enter the string to use in place of the string removed. No spaces are allowed. To insert empty spaces, use ` `. It only

makes sense to enter something else here if you have chosen 'replcont', 'replabort', 'replskip', or 'replendskip' as the action. Click 'Add' to generate the rule and add it to the list field.

- To edit a filter rule, select it from the list field by clicking it. The values are then displayed in the top half of the window. Modify the values and confirm the changes by clicking 'Edit'.
- To delete filter rules, select them from the list field and click 'Delete'.

MIME Type Filter

Content filtering on the SuSE Firewall on CD is carried out by the httpf proxy. Configure it in Figure 3.37 on the facing page. Part of the content filtering is triggered by the tags contained in an HTML page.

Content filtering can take place with the MIME types of documents accessed from the web over HTTP. These MIME types include, for example, images (GIF, JPEG, ...), audio files, (MPEG, WAV, ...) or videos (MPEG, AVI, ...). By means of the bit stream transmitted, it can be determined to which MIME type the requested document belongs. This makes it possible to detect specific file types and allow or refuse transmission.

If httpf detects a MIME type that is not defined, the object in question is refused. This also makes it possible to determine if the MIME type transmitted matches the contents transmitted or if, for example, an executable program is pretending to be a gif image.

Operation

The dialog 'MIME Type Filter' shows a list of all MIME types already configured. The three fields 'MIME type', 'Offset', and 'String' are for creating a new type or editing a Mime type. 'MIME type' specifies which MIME type is involved. Offset refers to the point after the beginning of the file at which the "string" is located. The two parameters 'Offset' and 'String' are optional, because not all file formats can be uniquely identified.

All MIME types intended to pass through the filter must be defined in this interface. A configuration with most MIME types is contained in the default template. If you do not want to let certain MIME types through, mark them and click 'Delete'. With 'Next', continue to the next dialog.

Parent Proxy Configuration

If your provider makes available a specific proxy for HTTP requests, the IP address of the proxy can be entered in 'IP address of the parent proxies'.

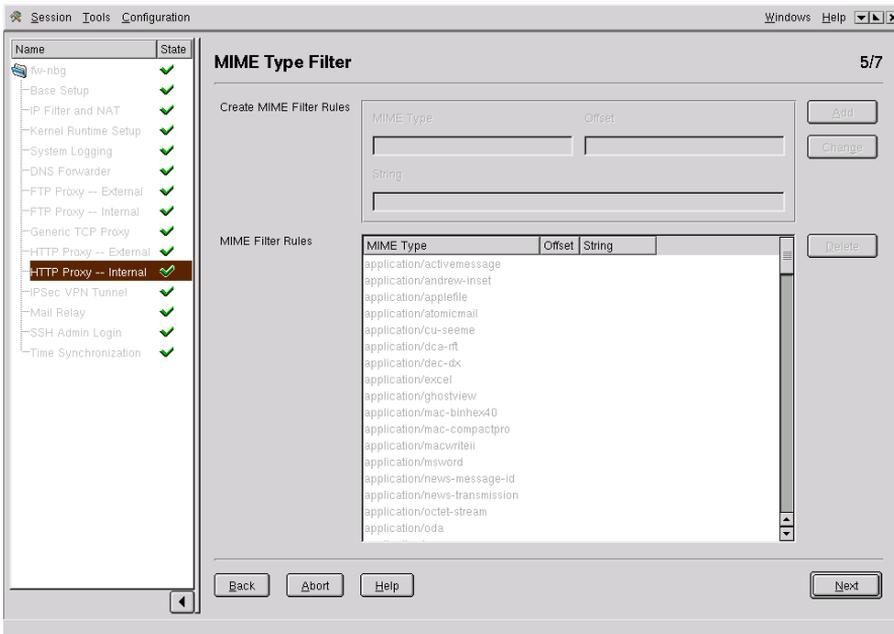


Figure 3.37: Mime Type Filter

Enter the appropriate HTTP port of the provider is entered in 'Parent Proxy Port'. Refer to Figure 3.38 on the next page. Confirm the settings with 'Next'.

Access Configuration

At the end of this module you can again activate the automatic configuration of IP filter rules and switch on logging of access violations. As usual, select IP addresses from the list for which access should be granted (see Figure 3.39 on page 83).

The Example, Inc., Configuration

Page 1/7 HTTP Proxy — Internal

The HTTP proxy should be operated on the port specified:

```
HTTP Proxy Port: 3128
Auto listen to:  activated
```

The proxy should be transparent for the internal network and, to spare bandwidth, the cache for the proxy should be activated.

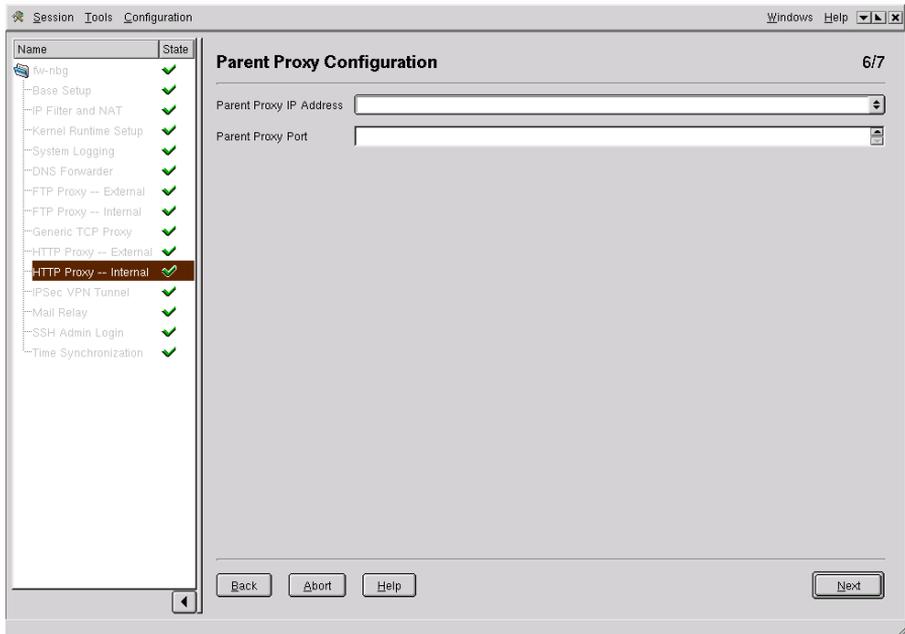


Figure 3.38: Configuration of the Internal HTTP Proxy

Proxy Mode activated
 Cache activated
 MB Cache 1000

Tip

Transparent Proxy and Caching

Because a transparent proxy is set up, there is no need to configure the proxy on the respective clients. Through caching, clients have the advantage that pages already in the cache are displayed more quickly.

Tip

Page 2/7 HTTP Proxy — Internal

At this point, ACLs (Access Control Lists) are defined. It is possible with these to precisely specify who may see which pages. In Example, Inc., access is simply granted to the entire internal network.

Name of ACL Nuremberg internal
 Type src

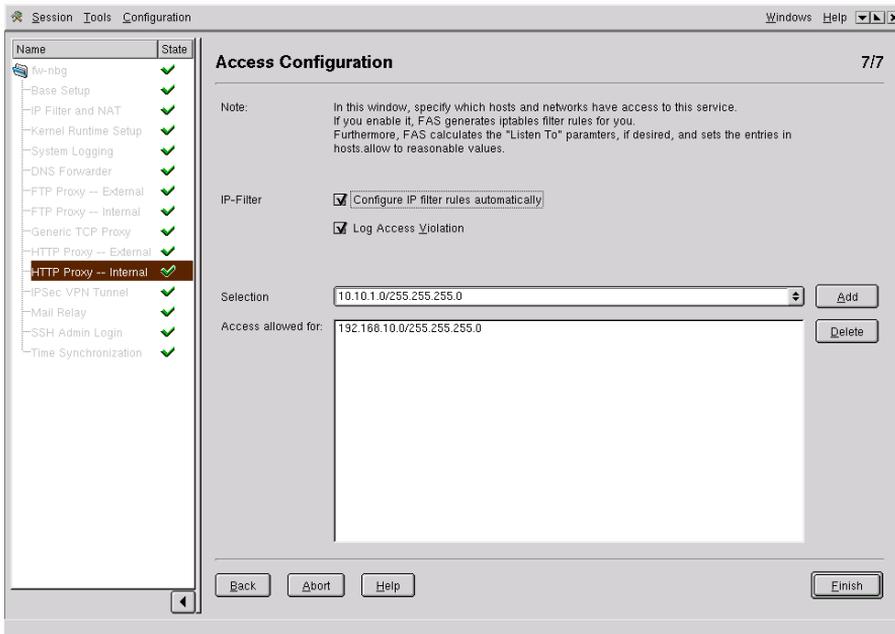


Figure 3.39: Access to the Internal HTTP Proxy

The values for the ACL are entered by selecting them then editing them. As the value, the network `192.168.10.0/24` should be entered.

Page 3/7 HTTP Proxy — Internal

The ACL generated is added to the configuration here. Mark ACL → Nuremberg internal and insert it with Add. The rules are processed from top to bottom. The order here is important. To move the rule to its correct place, select it then click (↓) until it is directly above deny all.

Page 4, 5, and 6 of HTTP Proxy — Internal

At Example, Inc., no use is made of the content filter. This is why the MIME type filter cannot function. No parent proxy is made available by the provider, so this is also not configured.

Page 7/7 HTTP Proxy — Internal

The proxy should only be used by clients from the internal network. For this reason, access is restricted to this network:

```
IP Filter          activated
Access allowed for: 192.168.10.0/24
```

IPsec VPN Tunnel

Use the VPN connection module to set up VPN networks. These virtual private networks can be regarded as a tunnel between two hosts that runs through the Internet. This tunnel knows nothing about the information transmitted in it.

The VPN networks are implemented on the SuSE Firewall on CD with IPsec. IPsec is a protocol family enabling secure connections to be established between computers. Authentication is by means of certificates. An additional possibility for authentication is shared keys.

Data sent through this tunnel is automatically encrypted. The certificates required for authentication are generated or imported with the certificate management of FAS, explained in [Certificate Management](#) on page 105).

Selecting the Local Certificate

In the first dialog of the module, shown in Figure 3.40 on the facing page, select an X.509 certificate for authentication. This certificate is used on the firewall host for which this configuration will be used. To create certificates, see [Certificate Management](#) on page 105.

If you do not want to use strong authentication, you do not have to use a certificate. See Figure 3.41 on page 86.

Note

“Strong authentication” means authentication with a key and a pass phrase.

Note

VPN Connections

In the second dialog, set up individual VPN connections. All VPNs configured until now are displayed in a table.

Select ‘Add’ to set up a new VPN connection. In the dialog that appears, see the tabs ‘General Settings’, ‘VPN Connection’, ‘Authentication’, ‘IP Filter’, ‘Masquerading’, and ‘Destination NAT’.

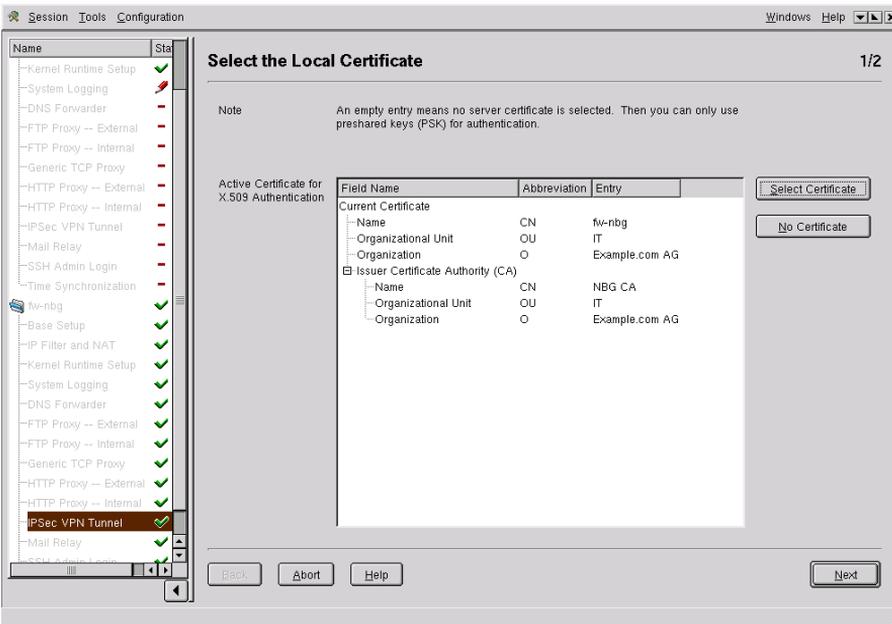


Figure 3.40: Select the Local Certificate

General Settings

This dialog is shown in Figure 3.42 on page 87. Assign a name for the connection. This makes it easier to identify the connection if you set up several tunnels. A connection name may only contain letters, digits, underscores, and hyphens.

Decide if the SuSE Firewall on CD should open a connection to another VPN server when it boots (Client mode) or it should wait for incoming VPN connections (Server Mode). For external locations, which should have a connection to the headquarters and have real IP addresses, the client mode should be chosen both for the external location and the headquarters.

For 'PFS Setting', "Perfect Forward Secrecy" is activated by default. This means that the asymmetric key may be changed. Next, specify the lifetime of the key in minutes. After this time, an attempt is made to change the symmetrical key. Specify the number of attempts to make in case of error. If you enter 0, there is no limit to the number of attempts.

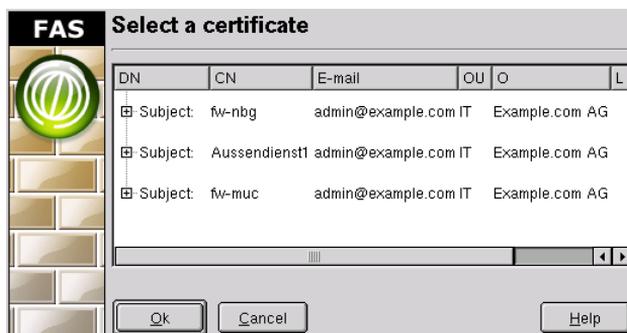


Figure 3.41: Available Certificates

VPN Connection

The 'VPN connection' tab is shown in Figure 3.43 on page 88. Under 'Local Configuration', activate the check box 'Route a subnet' if a local subnet should be routed. If the check box is activated, specify the subnet to contact with an incoming VPN connection. For 'Subnet', the network must be in a different IP address range than that on the remote side.

Under 'Remote Configuration', choose between dynamic IP addresses ('Road Warrior') and a fixed IP address ('Fixed IP Address'). The Road Warrior configuration enables a client to establish a connection to the VPN server from any IP address (e.g., dialing into the Internet from any provider, making available access to the company network).

If you decide on a fixed IP address, specify this in the next entry field. If a remote subnet should be routed, activate 'Route a Subnet' under 'Remote Subnet'. In this case, enter the address of the subnet in the next entry field. 'Subnet' refers to the network segment that should be accessed by the tunnel. The subnet must be in a different IP address range than the one on the local side. When setting up the tunnel, a route is automatically set to this subnet.

Authentication

Determine the authentication mechanism for the connection (see Figure 3.44 on page 89). Use X.509 certificates (recommended) or authentication using shared keys. The X.509 certificate is comparable to an identity card issued for a computer. With this certificate, the computer is authenticated with all VPN locations. Only one certificate can be selected per computer.

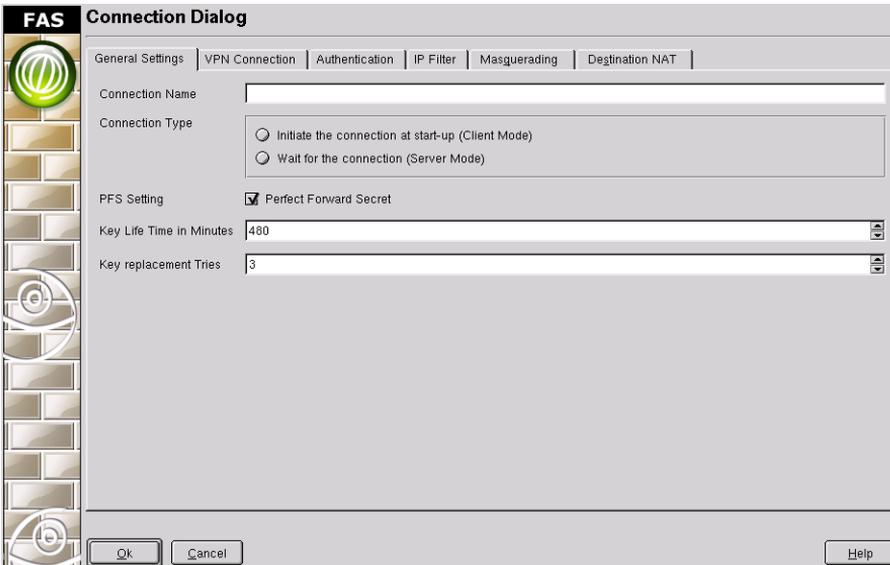


Figure 3.42: VPN: General Settings

A shared key is a random string. Quotation marks (") may not occur in the shared key. Every connection can have its own shared key, but it must be the same at both ends of a tunnel.

Transmitting the shared key must take place in a secure manner, because anyone who knows this key can authenticate himself at the VPN gateway and obtain access. For this reason, it is preferable to use X.509 certificates. If you use dynamic IP addresses, you cannot use shared keys.

If you activate authentication via certificates, select the corresponding certificate with 'Select'. To use a shared key, enter this in the corresponding entry field.

IP Filter

In this dialog (see Figure 3.45 on page 90), allow or deny packets arriving through the VPN tunnel on the basis of the defined filter rules. The check box 'allow all' allows everything. If this is not activated, you can only allow certain packets for the VPN tunnel through the definition of filter rules. In the first half of the mask, define the rules. The second half contains an overview window displaying all defined rules.

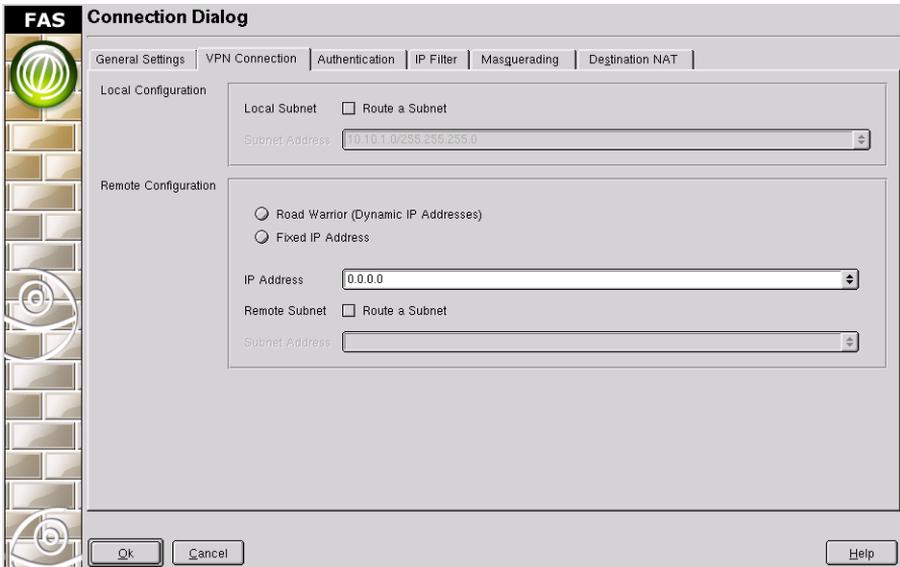


Figure 3.43: VPN Connection

To define a rule, first select the transmission protocol: `tcp`, `udp`, or `icmp`. If you select `tcp` or `udp`, you only need to specify the destination port or port range. If you enter `icmp`, choose the message type from the list.

Activate 'Log Access Violation', if desired.

Masquerading

In this dialog, activate masquerading for the collection, as shown in Figure 3.46 on page 91. This is only possible, however, if a local subnet is routed, but not a remote one. All packets that leave the tunnel on the local side are then masqueraded.

Destination NAT

In the final dialog for VPN configuration, shown in Figure 3.47 on page 92, configure the rule for "Destination NAT". Select the transmission protocol (`tcp` or `udp`), enter the destination port (a port or port range), then enter the IP address and the port to which the packets should be redirected.

If you click 'Add', the rule appears in the overview window. With 'Change', edit an existing rule. 'Delete' removes a rule.

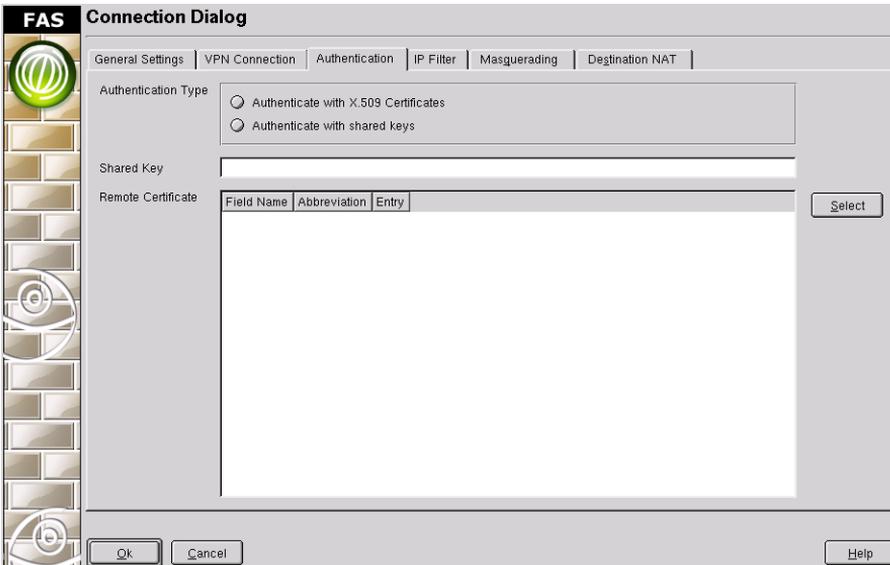


Figure 3.44: Authentication for a VPN Connection

When satisfied with all the dialogs, click 'OK'.

Information about setting up VPN on Windows clients can be found in *IPsec Client on Windows XP and Windows 2000* on page 131.

The Example, Inc., Configuration

To configure the IPsec module, the required certificates must first be generated. This is done in 'Modules' → 'Certificate Management'. First, a ROOT CA must be generated. This can only be done once and is only needed once. All further certificates are signed with this certificate.

Select 'Certificate Management' and click "Create CA". All fields must be completed. Here is an example certificate:

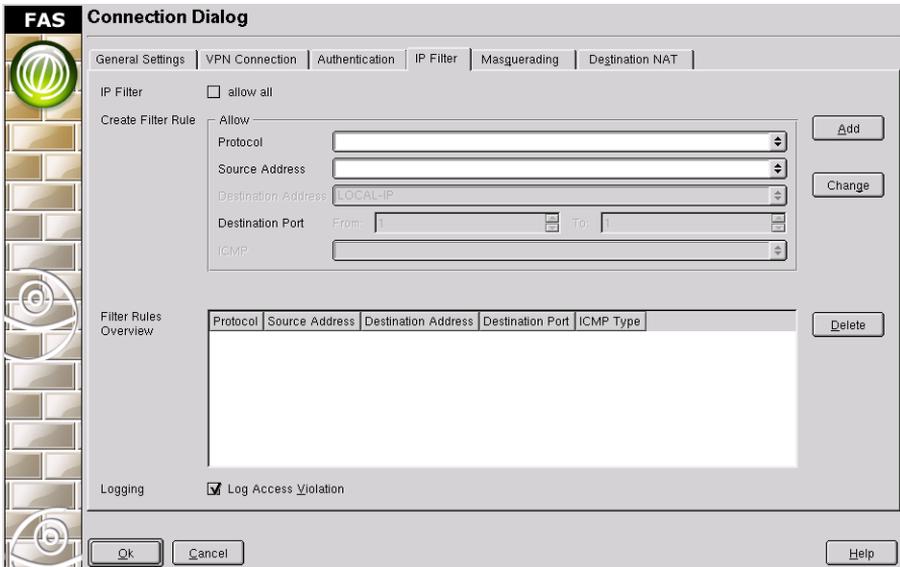


Figure 3.45: Filter Rules for a VPN Connection

Common Name	RootCA
E-mail Address	admin@example.com
Organizational unit	edv
Organization	Example, Inc.
Locality	Nuremberg
State	Bavaria
Country	DE
Days	1825
CA Password	PvdFSiN2
Verify CA Password	PvdFSiN2
Key size	2048

The certificate for the Firewall itself is created in the 'Certificate Management' → 'Create Certificate':



Figure 3.46: Masquerading for a VPN Connection

Common Name	Firewall-nbg
E-mail Address	admin@example.com
Organizational unit	edv
Organization	Example, Inc.
Locality	Nuremberg
State	Bavaria
Country	DE
CA Password	PvdFSiN2
Certificate Password	Firewall
Verify Certificate Password	Firewall
Key size	2048

Further certificates are required for the other locations so are generated according to the same scheme:

Firewall-fam	certificate for Frankfurt
Firewall-muc	certificate for Munich
sales staff1	certificate for sales representatives
to sales staff10	

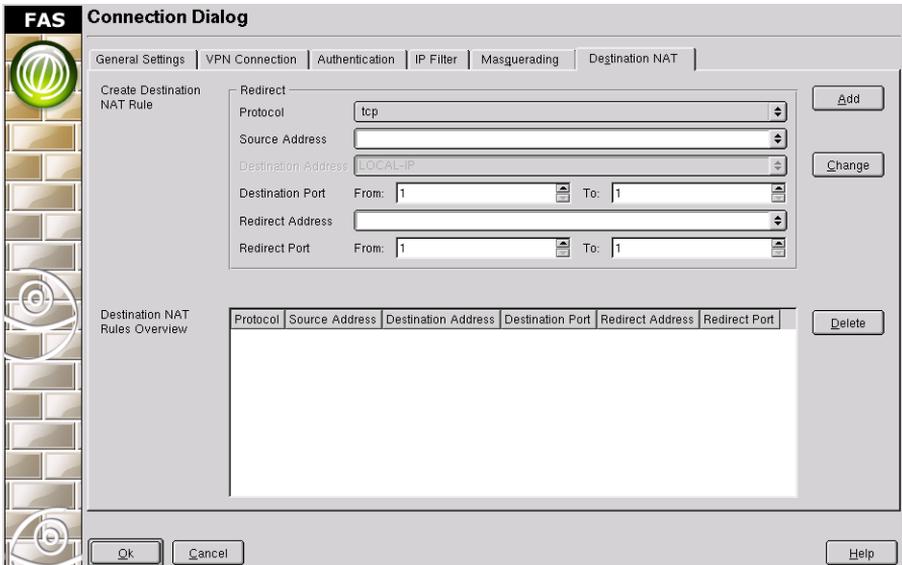


Figure 3.47: Destination NAT for VPN Connections

Page 1/2 IPsec VPN Tunnel

First, the certificate for Nuremberg is selected. To do this, click 'Select Certificate' and select the certificate `Firewall-nbg`. The password in the example configuration is `Firewall`.

Page 2/2 IPsec VPN Tunnel

Now the connections to the other branches and to the sales representatives are configured.

1. General Settings

The firewall in Nuremberg is the master. For this reason, it is set as follows:

Connection Name `Firewall-Frankfurt`
 Connection Type `Client`

The other settings for "PFS Setting", "Key Life", and "Key replacement Tries" remain unchanged.

2. VPN Connection

The Frankfurt branch should have full access to the internal network:

```
Local Subnet    activated
Subnet Address  192.168.10.0/24
```

The location in Frankfurt has the IP 100.100.100.2. The Frankfurt subnet should also be available in Nuremberg:

```
Fixed IP Address  activated
IP Address        100.100.100.2
Remote Subnet     activated
Subnet Address    192.168.11.0/24
```

3. Authentication

Authentication should take place with the certificates. For this reason, the field “Authenticate with X.509 Certificates” is activated. With ‘Select’, the Frankfurt certificate “Firewall-fam” must be selected.

4. IP Filter

The branch in Frankfurt is considered trustworthy. For this reason, all traffic through the tunnel is allowed.

5. Masquerading

Since no other internal subnets should have access to the VPN tunnel, nothing needs to be configured here.

6. Destination NAT

Every computer may communicate with every other computer. This is why no rules are required here.

In Munich, almost exactly the same configuration is required. For the sake of clarity, here are the exact modifications:

1. General settings

```
Connection Name  Firewall-Munich
```

2. VPN Connection

```
IP Address        120.120.120.1
Subnet Address    192.168.12.0/24
```

3. Authentication

The certificate “Firewall-muc” must be selected.

To conclude the VPN configuration, here is an example of the configuration for the first sales representative. This is somewhat more complicated:

1. General settings

Connection Name SalesReps

Connection Type Wait for the connection (Server Mode)

Settings for PFS Setting, Key Life Time, and Key Replacement Tries remain unchanged.

2. VPN Connection

The sales representative should have access to the internal network. The settings for this are as follows:

Local Subnet activate Route a Subnet

Subnet Address 192.168.10.0/24

All dial-up takes place with dynamic IP addresses. Therefore, Road Warrior is activated.

3. Authentication

Certificates are used to check access permission. Therefore, Authenticate with X.509 Certificates is activated.

4. IP Filter

No restrictions should be placed on sales representatives. IP Filter → allow all is activated.

5. Masquerading

No configuration is necessary here.

6. Destination NAT

Since the sales representatives may address all computers in the various subnets, no rules are required here.

Configuring the Mail Relay

To use this module, you must have a hard disk mounted in the base configuration and the directory `/var` must be located on this hard disk. The module is shown in Figure 3.48 on the facing page.

Internal Mail Server Enter the IP address of the name of your internal mail server.

Forward following domains Enter the domain names that should be accepted by the mail relay and forwarded to the mail server.

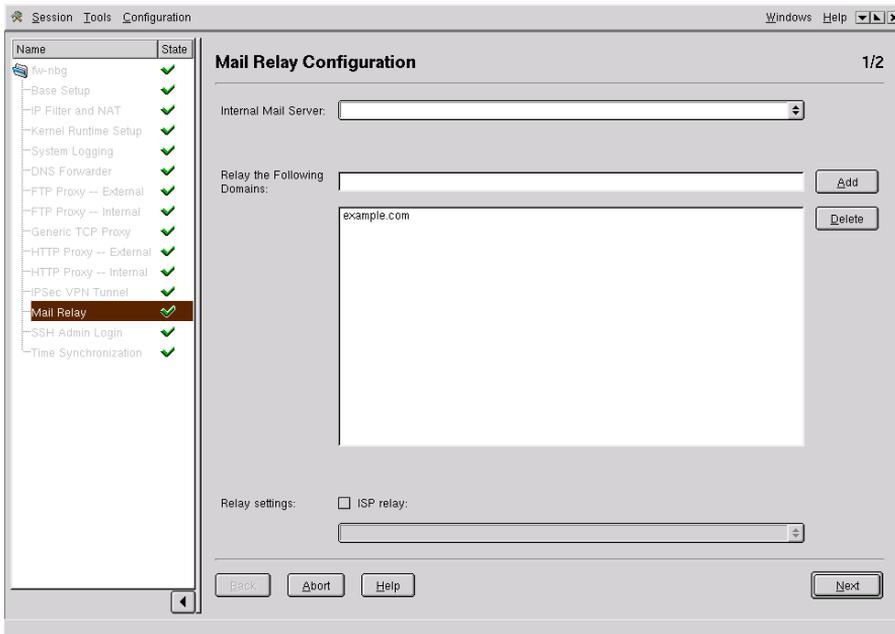


Figure 3.48: Configuration of Mail Relay — Dialog 1

ISP Relay Activate this check box to forward all outgoing e-mails to the SMTP server of your Internet service provider. Enter the IP address of the SMTP server of your provider.

With 'Next', continue the second part of the mail relay configuration. This is shown in Figure 3.49 on the next page. The following details are required here:

Local networks Enter a list of networks that may send e-mails using the firewall (e. g., 192.168.0.0/24).

Auto listen to Activate the automatic calculation of the listen-to parameters or manually select the corresponding IP addresses.

Listen to selected IP addresses Enter the IP addresses on which the mail relay accepts requests, if you have decided against automatic calculation.

IP Filter Activate this if the IP filter rules should be generated automatically. If required, activate the logging of access violations.

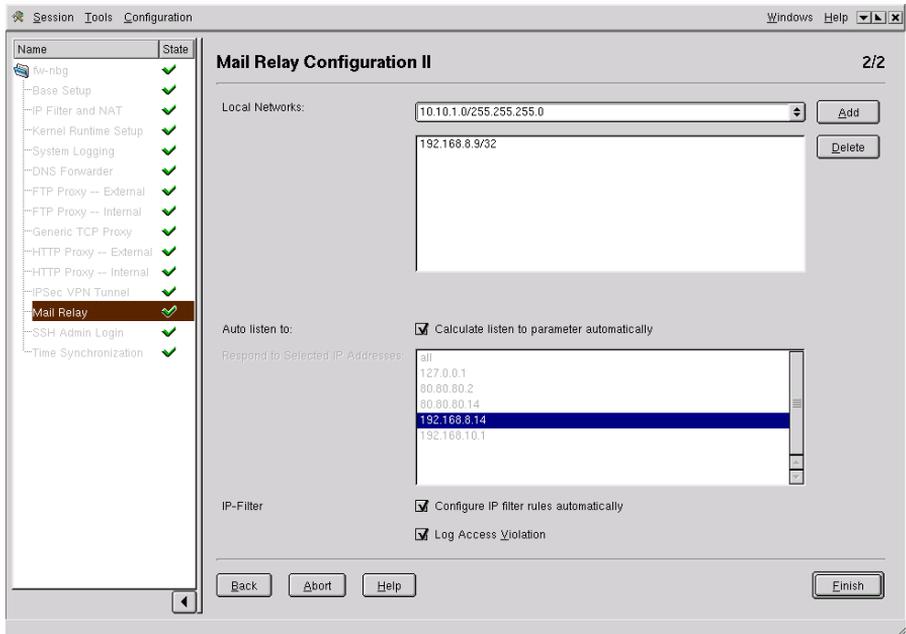


Figure 3.49: Configuration of Mail Relay — Dialog 2

With 'Next', complete the configuration of the mail relay.

The Example, Inc., Configuration

Page 1/2 Mail Relay

The mail server of Example, Inc., listens to the IP address 192.168.8.10. It must be determined for which domains it should accept e-mails and to where it should send outgoing mails:

Internal Mail Server	192.168.8.10
Relay the Following Domains:	example.com
ISP relay:	activated
	80.80.80.70

To avoid DNS problems here, only IP addresses may be given.

Page 2/2 Mail Relay

All mail traffic should pass across the internal mail server. For this reason, the mail server is also the only computer that may send mail via the mail

proxy. With the netmask, this is reduced to a single computer:

```
Local Networks 192.168.8.10/32
```

All other parameters for Example, Inc., remain in their default settings.

Administration via SSH

In this module, configure SSH access for the firewall administrator. For the encrypted connection to the firewall machine, the corresponding SSH keys are entered here. More detailed information about SSH can be found in *SSH — Secure Shell, the Safe Alternative* on page 190.

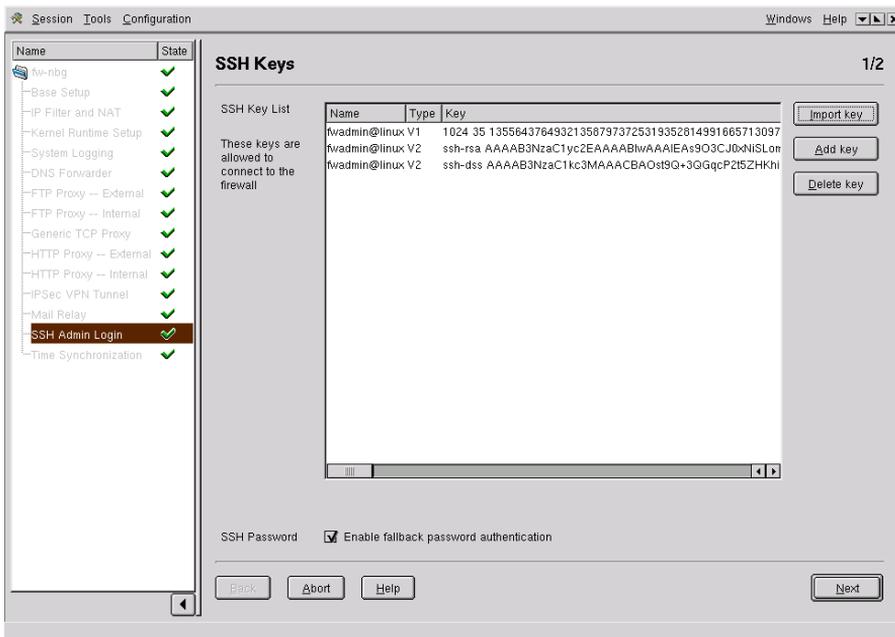


Figure 3.50: Add or Delete SSH Keys

In the first mask of this dialog, shown in Figure 3.50, save your “ssh-public-key” for access as `root` on the firewall, so you can later authenticate yourself and log in to the firewall machine. You have the following options:

- Import a key. An “ssh-key” normally lies in the user’s home directory in the directory `.ssh/` in the file `identity.pub`. If you are using ver-

sion 2 of SSH, the files `id-dss.pub` and `id-rsa.pub` are involved. Select them and the key appears in the list (see Figure 3.51).

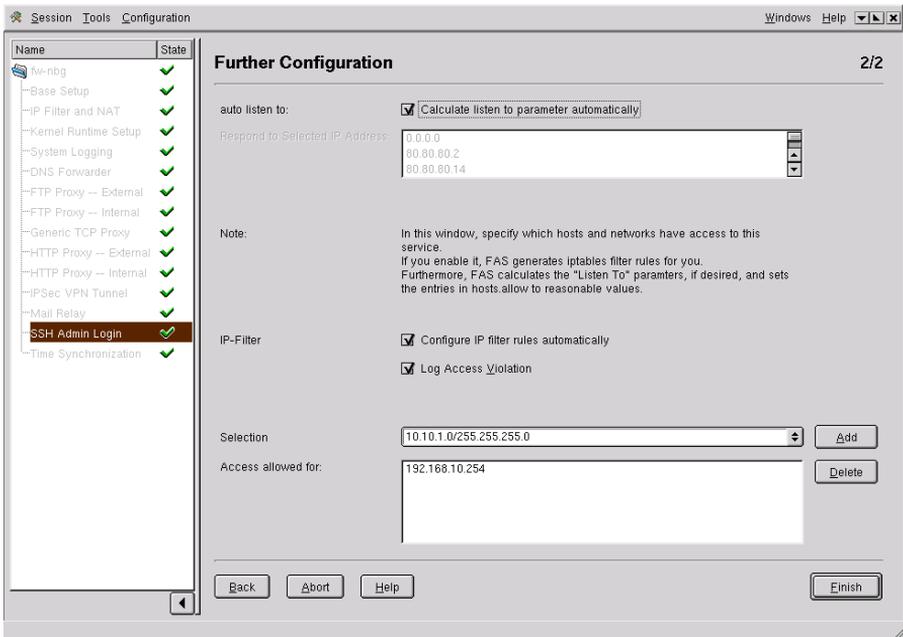


Figure 3.51: Import Key

- Enter a key by copying and pasting. If you click ‘Add Key’, a dialog opens. In the text field, enter your “ssh-public-key” and confirm with ‘Ok’. The scheme now appears in the lower list.
- With ‘Delete Key’, remove a selected key.
- Allow password authentication in emergencies by activating the corresponding check box. We strongly advise only allowing access via RSA keys.

On the second dialog page, assign access rights for SSH. Have the “listen-to” parameters automatically calculated by activating this check box or select the desired IP addresses yourself.

If required, activate the automatic generation of IP filter rules and the logging of access violations. At the end of the input mask, select the IP addresses from which SSH access should be enabled.

The Example, Inc., Configuration

Page 1/2 Administration via SSH

Remote access to the firewall should be possible only by the Adminhost via SSH. To enable this, the firewall must have the public key of the administrator. To do this, select 'Import key' and choose either `id_rsa.pub` (SSH2) or `identity.pub` (SSH1). If required, both SSH versions can be supported. To be extremely cautious, deactivate the fallback to "password authentication" by unchecking `SSH password`.

Page 2/2 Administration via SSH

At Example, Inc., the Adminhost has the IP address `192.168.10.254`. This must be entered here:

```
Access allowed for: 192.168.10.254
```

Time Synchronization

With this module, configure the time server `xntpd`. This ensures that computer time can be kept in sync with an external time source (a computer with the exact time). This is important so the time stamps in log files can be compared to the time stamps in other log files on other hosts.

Enter the IP addresses of the time servers from to request the current time (Figure 3.52 on the next page). The NTP protocol is UDP-based, which means you must open the corresponding ports of the packet filter in FAS.

The Example, Inc., Configuration

So all computers in the internal network are always set to the same time, the NTP service is used. The firewall itself is synchronized from the time server. At Example, Inc., there is the following time server:

```
Time Server: 192.168.10.23
             192.168.10.24
```

Log File Analysis

The Log File Analysis in FAS is opened with 'Tools' → 'Log File Analysis'. End the log analysis by clicking 'Log File Analysis' → 'Finish'. Get context-sensitive help with 'Log File Analysis' → 'Help'.

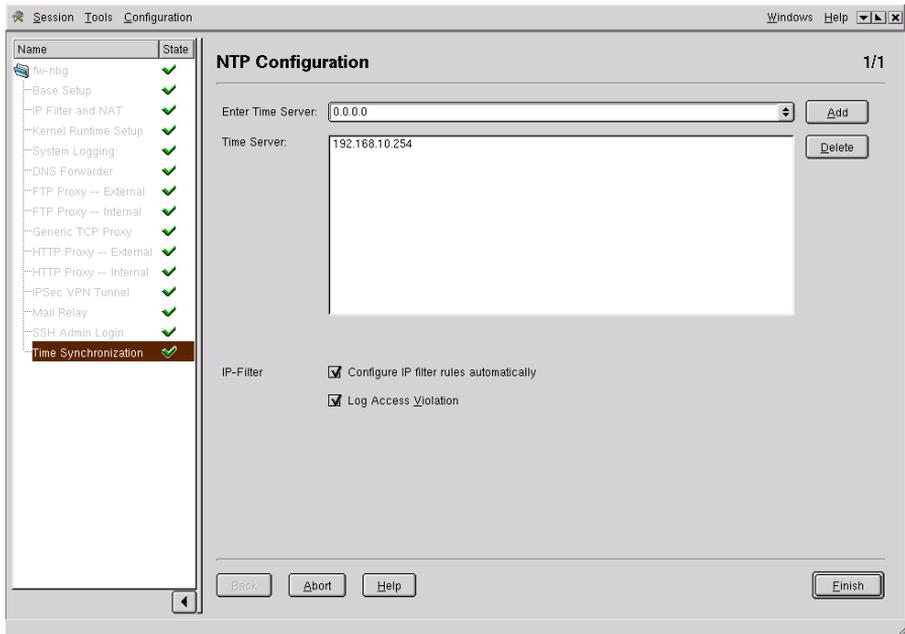


Figure 3.52: Specifying the NTP Time Server

Assuming you have activated the forwarding of log files for all active firewall configurations to the Adminhost, this module provides convenient access to all recorded accesses and sender statistics of the networks monitored. In the left-hand side of the window, see an overview of all active configurations whose log files have been transferred to the Adminhost.

Click the directory of the configuration for which to view the log analysis. A tree view appears of four groups of log files: 'IP Filter Statistics', 'View Interface Statistics', 'View Log Files', and 'View Mail Statistics'. In the right-hand side of the window, see information about the currently active configuration:

SYSLOG Configuration Name of the currently selected configuration

Host Name Name of the system on which the SuSE Firewall on CD is operated with this configuration

User Comments comments on the configuration

If you click a name, an information page appears in the right-hand side of

the window, providing a summary of the data recorded. This view is divided as follows:

SYSLOG Configuration Name of the current configuration

Description A short description of the data recorded

Available Statistics and Log Files In a table, see which files can be analyzed and what type of data are contained there.

The information page provides a rough overview of the available data. By clicking the desired statistics in the left-hand side of the window, see the data either in graphic form or as a list that can be searched.

The Log Files

At this point, all the log files of your firewall that are recorded and transmitted through syslog are bundled together. See the file name, file path, the date of first creation, and the date of the last modification in an overview table. The log files managed are:

ipfilter Kernel messages of violations against IP filter rules

kernel All other messages from the kernel

local Messages logged with the syslog facility 'local[0-7]'

warn Warnings from individual programs and from the kernel (logged via the syslog facility warn)

mail Messages from the mail subsystem

vpn_updown Status of the VPN module. This log file only appears if the VPN module is installed.

messages All other messages from syslog.

All statistics recorded can be searched from a search mask after clicking the corresponding name in the tree view.

Evaluating the Log Files

Log files often become very large, depending on the type of data recorded and the duration of recording. With the search mask of the Log File Analyzer, restrict the extent of the data to analyze. Fill out the mask then click 'Show' to display data in the lower part of the window filtered according to the search criteria. Using 'First Page', 'Previous Page', 'Next Page', and 'Last Page', browse through the pages of the display.

With the following instructions, determine the extent and the compilation of the data displayed:

Max Lines per Page Maximum number of lines per display page

Begin Date Precisely define the beginning date of the data to display.

End Date Defines the end date.

Regular Expression Search for keywords through a regular expression. If you are interested, for example, only in messages containing the string "ICMP", enter ICMP here.

The IP Filter Statistics

The IP Filter statistics for the firewall can be evaluated in three different ways:

alldates All data accumulated since the beginning of the recording date is displayed graphically (see Figure 3.53 on the facing page).

lastweek All data of the past week is displayed graphically.

today All data for the current day is displayed graphically.

Click the statistics required in the left window. On the right-hand side, an HTML page appears with the detailed evaluation.

General Summary The general summary includes information about the number of blocked packets, the average blocking rate per day, the number of individually blocked packets, the number of hosts whose packets were blocked, the entire volume of blocked packets, and the average size of them.

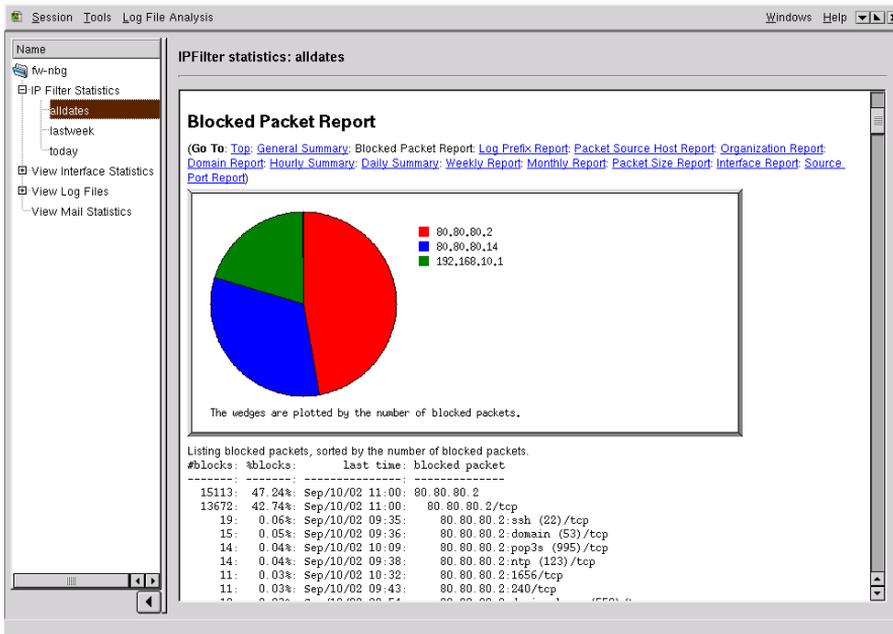


Figure 3.53: IP Filter Statistics

Blocked Packet Report A pie chart shows the proportion of packets blocked grouped by source IPs. The blocked packets are also shown in table form, sorted according to the number of blocked packets.

Log Prefix Report As with the 'Block Packet Report', the report is a pie chart in which the number of blocked packets for individual log prefixes is shown. The table is sorted by the number of packets blocked.

Packet Source Host Report In the pie chart, the source hosts (resolved name and domain) are allocated to the blocked packet amounts which were sent by them. In the overview in table form the packet amounts and sizes are listed for all hosts from which more than 0,5 percent of the blocked packets originate.

Organization Report In the pie chart, the number of blocked packets are shown by source organization. The table shows the name of the source organization, the number of packets blocked, the percentage of the total amount of all packets blocked, and the percentage of the total blocked packet volume.

Domain Report This report is similar to the 'Organization Report', however, sorting is done according to domain extensions, such as .com.

Packet Size Report In an overview in table form, the size of blocked packets is listed with their number and the percentage of the overall amount.

Interface Report The amount of packets arriving and blocked packets are displayed by the name of the interface. The overview in table form shows the interfaces (with at least five blocked packets), the number of blocked packets, and the percentage of the total volume of blocked packets.

Source Port Report In the overview in table form, the twenty source ports from which the most packets were blocked are listed. They are assigned the number, the size of the blocked packets, and the percentage value.

The Interface Statistics

The interface statistics contain, for each interface of the firewall, a detailed substatistic displayed according to packets sent and received and to data amounts. It also shows any errors that have occurred. A sample is shown in Figure 3.54 on the next page.

For each interface, it shows two graphs in four different time resolutions (daily, weekly, monthly, and yearly). The sent and received data amounts or packet numbers per minute are shown on the Y axis. A specific time interval is shown on the X axis.

Mail Statistics

The mail statistics consist of three graphs, illustrating different aspects of the daily mail traffic on your firewall:

Day Graph the amount of received and sent mails per minute by time of day.

Day Graph of Errors The amount of errors occurring per hour. "Errors" include bounced mails and rejected mails.

Day Graph of Mail Traffic in Bytes The data amounts of received and sent mail (in KB) per minute, by time of day.



Figure 3.54: Interface Statistics

Certificate Management

Access the certificate management module in FAS with 'Tools' → 'Certificate Management'. Close certificate management with 'Finish' from the same menu. Certificates for encryption when using IPSec with X.509 certificates can be generated, imported, and managed. In the main window, shown in Figure 3.55 on the following page, already existing certificates listed.

The keys and certificates are generated on the SuSE Adminhost for Firewall with the program package OpenSSL. With SSL, asymmetrical encryption is performed, which means one key pair, consisting of a public and a private key, is always necessary for encryption and decryption.

With the asymmetrical PKI encryption, the public keys are exchanged between the client and server. Encryption takes place using the public key of the recipient. Decryption requires the private key.

To sign a certificate, a CA (Certificate Authority) is needed. There are some locations with official CAs. You must pass all your certificates to this point to

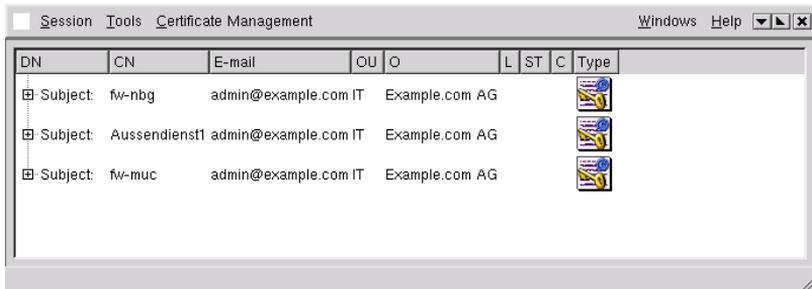


Figure 3.55: List of Certificates

have them signed. You can also generate your own CA and sign your certificates yourself. This is sufficient for most purposes.

Creating a Certificate Authority

You can only run this dialog once. The CA is globally valid for all configurations generated on this Adminhost. All certificates created are signed with it.

Select 'Certificate Management' → 'Create CA'. A dialog appears, which you should fill out completely. Many of these settings are used to generate client certificates.

- First, you need a name for your CA, such as your company name.
- Under E-mail address, enter a member of staff responsible for the CA (for example, the security manager for your company network).
- Department: Enter the department name here.
- Company: Enter your company name.
- Location: The location, such as that of your company headquarters.
- State: state
- Country: the two letter country code for your country (DE, US)
- Days: length of validity of the CA in days
- CA password: enter a password for the CA to prevent it being mis-used. Repeat it to confirm the password.

- **Key size:** Choose the key size here. A longer key is more difficult to hack. Choose 1024 or 2048 bits.

After you have made all necessary settings, confirm with 'Ok'. If the creation of the Certificate Authority (CA) runs successfully, you will see a message to that effect. Confirm by clicking on 'Ok'.

Creating a Certificate

Now you can create certificates signed with your CA. Select 'Certificate Management' → 'Create Certificate'. A dialog window like Figure 3.56 on the next page appears in which to enter the necessary details for a new certificate.

Some details have already been taken from the CA, so you only need to choose a new name for the certificate (e. g., the name of the computer for which the certificate is created), enter the previously defined CA password, and enter a new certificate password, which is repeated for confirmation. Choose the size of the key and click 'Ok'.

After successfully generating the certificate, see a message of completion, which you must confirm with 'Ok'. Your newly created certificate now appears in the list.

Deleting a Certificate

Choose a certificate to revoke from the list. Select 'Certificate Management' → 'Revoke certificate'. Enter the password for the CA with which the certificate was generated.

Confirm the deletion with 'Yes'. An updated list of certificates is displayed.

Importing Certificates

Select 'Certificate Management' → 'Import Certificate'. Choose the file containing the certificate from the directory listing that appears. A certificate to import can exist in the following formats: DER, PEM, and PKCS12.

No further details are required for the formats DER and PEM. If the certificate is in the PKCS12 format, a password dialog appears. Enter the import password set while exporting. Then set a new password for the certificate. Confirm it by entering it again.



Figure 3.56: Dialog for Creating Certificates

Exporting Certificates

Select 'Certificate Management' → 'Export Certificate'. There are three different formats in which a certificate can be exported: PEM, DER, and PKCS12.

To save the certificate in the PEM or DER format, select the certificate to export then select 'Certificate Management' → 'Export Certificate'. In the dialog that appears, choose the location to which to save and enter a name for the file. Choose the format: DER or PEM. The certificate is then saved. When exporting in the PEM format, also specify if the certificate should be saved on its own or with keys.

The certificate can only be exported in PKCS12 format if a key exists for this certificate. Proceed as for PEM and DER certificates. Choose a name for the certificate. For the file format, select PKCS12. Enter the password for the certificate then enter an export password. Repeat this for confirmation purposes. You must enter this password to import the certificate.

Saving the Configuration

Save your configuration by clicking 'Configuration' → 'Save' or 'Save All'. If you try to leave the configuration but have made modifications, you will be asked if you want to save them.

The configuration is saved on the hard disk under

```
/var/lib/fas/<username>/configs/<configurationname>/floppy/
```

No normal user on the Adminhost can read this configuration. Only root has the permissions to do this.

To create a floppy, insert one into the floppy drive, select the configuration to save, then select 'Configuration' → 'Write Floppy'. Now the floppy is written. Before this is possible, the configuration must be recognized as properly configured, which is shown by a green check mark in the left-hand side of the dialog.

Editing an Existing Firewall Configuration

To edit an existing firewall configuration, start the FAS. Log in as the user who created the configuration with 'FAS' → 'Login' and enter the corresponding password. FAS then lists all configurations created by this user.

Choose the configuration to edit from the list. Double-click the configuration name in the list in the left window to edit one of the services in the corresponding module.

Editing Configuration Files

Open the file editor with 'Configuration' 'Edit Files'. You can modify configuration files directly with this editor. Only do this if you really know what you are doing the effects your manual modifications will have. The FAS does not check these modifications for correctness.

On the left, the overview window shows the main directories etc, root, and sbin. The large window next to this displays the contents of the selected main directory are displayed. Click a configuration file in this window to open the contents of the file in the editor where you can modify them by hand. Click a directory to list the files in it.

Via the 'File Editor' menu, create a new file or directory. Files can be imported and entries be saved or deleted. With 'Help', access context-sensitive

online help. To leave the file editor, select 'Finish' from the menu. Save your modifications to configuration files by pressing 'Finish'.

Testing the Configuration

The configuration created with the administration program still needs to be tested before it is used. To do this, start the firewall is started without any connection to the Internet or intranet. Connect the firewall directly to the Adminhost with a crossover cable.

Tests for the packet filter can be simply carried out using a port scanner. For this purpose, the program `nmap` is installed on the Adminhost. If you run the port scan from the Adminhost, make sure its packet filter of the Adminhost is switched off. Log in as the user `root` and enter `SuSEfirewall stop` on a console. This ensures that all returning IP packets can be accepted by the Adminhost. Do not forget to activate the packet filter after tests have been completed with `SuSEfirewall start`.

After the port scan, the result of the scan and the log file on the firewall should be documented and saved. Check the function of the following services:

- Test the name server, for example, with `nslookup` followed by the name of the firewall. Observe the log file, for example, with `grep named /var/log/messages`. No messages may occur containing error.
- Test the mail relays by sending an e-mail and observing the log files `/var/log/mail` and `/var/log/messages`. Search log files for postfix (`grep postfix /var/log/mail`)
- Test the FTP proxies, for example, `telnet` to your firewall.
- Test the HTTP proxy by trying to reach the Internet from a client.
- Test `ssh` by logging in from a client to the firewall.

Always check all procedures by using the log files.

Documenting Configuration, Tests, and Results

It is very important to document the configuration, the tests conducted, and their results. Keep a record of what is allowed or denied by the configuration and how this is guaranteed. Using such documentation helps find and remedy possible configuration errors. The documentation is also required for auditing the firewall.

Monitoring the Firewall

A firewall without continual monitoring is only effective to a limited extent. A number of tools are available on the Adminhost for monitoring the firewall. The most important source for information is the log files, which, depending on the configuration, are written by the firewall to the hard disk or to the log host.

The following programs are available as tools for analyzing the log files:

- FAS Log Analysis module
- xlogmaster
- logsurfer

You can use the following network or packet sniffers to monitor your firewall:

- ntop
- tcpdump
- ethereal

With the following port scanners, it is possible to check the firewall for open ports and to check the packet filter configuration:

- nmap
- nessus

In addition to this, you can also use your own shell or perl scripts.

SuSE Live CD for Firewall

The Live CD for the SuSE Firewall on CD is the executable part of the firewall. It is a minimal SuSE Linux, designed with security criteria in mind. This affects the programs available and the kernel itself. The SuSE Firewall on CD is an “Application Level Gateway”, meaning that, for security reasons, there should be no routing of IP packets. Forwarding requests to services is handled by applications (proxies).

Hardware Requirements	114
Description	115
Services on the Firewall	115
The Configuration Disk	125
Boot Parameters	130

Using proxies and not forwarding IP packets are not enough to prevent undesired Internet IP packets from reaching the intranet or vice versa. This firewall functionality is adopted by the kernel packet filter configured by `iptables(8)`. This is where the Live CD comes into play. The operating system and all the applications are located on a CD — on a read-only file system. A RAM disk, where the Live CD is mounted, is generated when booting. The original system can simply be restored by rebooting the machine. Updating the SuSE Firewall on CD is also very simple because of this: just replace the old Live CD with the new Live CD and restart your firewall host.

Configuring the system and the services is done by a disk on which all the necessary configuration files are saved. This disk is mounted read-only when booting. For extra security, the disk should also be write-protected. The data on the configuration disk is copied to the RAM disk and the disk itself removed from the file system.

Hardware Requirements

- The SuSE Live CD for Firewall can run on any i586 or better x86 machine. At least a Pentium II is recommended.
- A minimum of 128 MB RAM is recommended. The firewall can still function with 64 MB, but proxies should not be used without more.
- A 3.5-inch floppy disk drive is needed for the configuration disk.
- The system must have a bootable CD-ROM drive.
- At least two network interfaces or an ISDN card and at least one network card are also needed.
- More precise system requirements are documented in detail in [System Requirements](#) on page 6

This chapter does not provide any instructions for configuring firewall services. Instead, it offers technical documentation for the well-versed administrator willing to grapple with the internal aspects of the system. Details relating to the configuration files of the services provided on the Live CD are also described. For configuration purposes, the Firewall Administration System (FAS) on the Adminhost should be used.

Description

The SuSE Linux Live CD for Firewall is a live file system CD from which all the applications run directly. Theoretically, the firewall host could be operated without a hard disk. However, a hard disk for the cache and spool directories is required by proxy services, such as Squid or postfix. A hard disk is also required if you want to save the syslogd log messages locally. To set up the firewall system, insert the CD and the configuration disk created on the Adminhost then boot the host.

Services on the Firewall

Application level gateways, or proxies, are located on the SuSE Live CD for Firewall for the most common and essential Internet protocols and other services:

DNS (Domain Name System) IP addresses can be converted into “Fully Qualified Domain Names” and vice versa with `bind8`, but only in the form of a cache-only name server, which is used as a forwarder for requests from the internal network. It is not intended to be configured as a DNS server.

SMTP postfix transports e-mail in an SMTP relay configuration.

HTTP and HTTPS — the WWW protocol Squid, `httpd`, `tinyproxy`, `mod_proxy`, and Apache.

FTP For file transfers from one host to another, the program `ftp-proxy` from the `proxy-suite` is used.

SSH Remote logins with encrypted transmission is handled by `openssh`. Authentication is performed with RSA key pairs. This service simplifies the administration of the firewall on a running system.

rinetd Generic tcp proxy `rinetd`.

ntp The time server `xntpd` synchronizes time on the firewall machine with other computers in the internal network. This means that the time stamps of the logs are synchronized.

ipsec As IPSec VPN software, FreeS/WAN is used.

i4l ISDN support

pppoed DSL support

fasfw FAS net filter script

syslogd Daemon for system logging

iptables Packet filter

cron and logrotate Rotation of local log files

To complete the range of services offered on the firewall, the following software packages are also available on the Live CD, but use them at your own risk.

cipe VPN tunnel software

pptpd MS VPN tunnel server (point-to-point tunneling protocol daemon)

cron Daemon use to run commands on a set schedule, normally used for log rotation.

ippl IP protocols logger (portscan logger)

scanlogd Portscan logger

ipvsadm Linux Virtual Server (to configure the load balancer contained in the kernel)

snmpd To display interface statistics

sockd Dante server, Socks v4/v5 server (proxy)

zebra, ospf6d, ospfd, bgpd, ripd, ripngd Routing software (e. g., for BGP and OSPF)

All these programs are Open Source software. On the SuSE Live CD for Firewall, all processes run in chroot environments. To increase the security of the system even more, compartment is used for the Linux kernel contained on the Live CD (2.4.18).

iptables

A typical iptables filter rule is very simple in theory. It normally consists of four parts:

1. a basic operation with which the rule is inserted. This is typically run with the command `/sbin/iptables -A`.
2. the identification of the packet
3. a description of the packet to treat
4. a description of what should happen to the packet once it is found

Inserting Rules

There is a range of options available for manipulating filter rules. The basic commands are:

- adding new rules to a chain: `iptables -A`
- inserting a rule at a specific point in a chain: `iptables -I`
- replacing a rule at a specific point in a chain: `iptables -R`
- deleting a rule within a chain: `iptables -D`

In these commands, a “chain” must be specified in each case. These will be explained in the next section. The syntax would then be, for example, `/sbin/iptables -A INPUT ...`

Be very careful with all options except “-A”. It is very easy to get into unexpected difficulties. If you want to delete all rules except the default policy, use `iptables -F`.

The Course of a Packet

With iptables the user has three different filter tables available: `filter`, `nat`, and `mangle`. Each of these tables describes different “chains.” A “chain” is described by a list of filter rules. Each of these rules says: “if a packet header matches my description, here are my instructions for this packet.” If the packet header does not match, the next rule is queried. If all rules have been queried and there are no matches, the “default policy” is applied.

The `filter` rules chain is described by three tables: `INPUT`, `OUTPUT`, and `FORWARD`. The `nat` rules chain is also described by three tables:

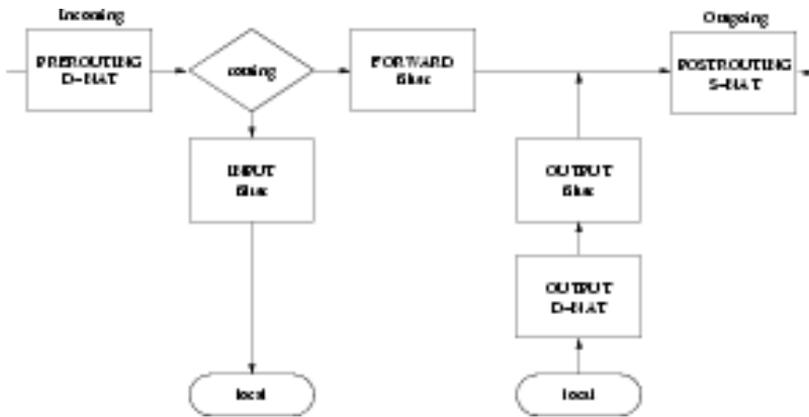


Figure 4.1: Course of a Packet with iptables

PREROUTING, OUTPUT, and POSTROUTING. Figure 4.1 attempts to illustrate the interplay of nat and filter tables.

Note

Course of a packet

All packets must pass through both a nat and a filter table before they can reach a computer program.

Note

The program iptables has an implicit parameter `-t [filter]` that, by default, is applied to the filter table. To address another table, specify it with the `-t [table name]` option. For example, to add a rule to the POSTROUTING chain in the nat table, enter the command `/sbin/iptables -t nat -A POSTROUTING...`

The implicit parameter `-t [filter]` is also valid for other operations. For example, the command `iptables -t nat -L -nv` lists the rules of the nat table. The commands `iptables -L -nv` and `iptables -t filter -L -nv` both list the rules of the filter table.

Packet Descriptions

With iptables, it is possible to check many features of packets. Only a few of these are mentioned here as examples. For more on this subject, see the man page for iptables (`man 8 iptables`). The following criteria are often used:

- Protocol type: TCP, UDP, ICMP
- Source port
- Destination port
- ICMP type
- Source address
- Destination address
- Interface: eth0, ppp0, etc.
- Inversion

Protocol type Protocols may be specified with their numerical protocol type or, for the special cases TCP, UDP, and ICMP, with their names. These names are not case-specific, so TCP is the same as tcp. A list of names together with their corresponding numbers can be found in `/etc/protocols`. The protocol specification can be negated by the use of a `!`. `!TCP` means all packets except TCP. If no protocol is specified, the rule is matched to all protocols.

Source and destination port In the case of “TCP” and “UDP”, a large range of additional options is available to restrict the selection of packets. Most frequently used is restricting the choice to a few source and destination ports.

There are two versions of port control:

1. TCP and UDP extensions
2. The multiport module

Both versions can be used to control the ports. There are a few significant differences, however. The UDP and TCP extensions always allow control over just one single port or a range of ports for a rule. Different source and destination port ranges may be specified.

The Multiport module allows up to 15 different ports to be specified for a rule. The order here is arbitrary. In contrast to the UDP/TCP extensions, however, only source ports or destination ports may be specified, but never both. Ranges of ports are also not allowed in this module.

ICMP Because ICMP packets do not use any port numbers, other selection criteria must be used. A list of possible parameters can be obtained with the command `iptables -p icmp --help`. Some frequently used types:

Name	Number
echo-reply	0
destination-unreachable	3
source-quench	4
echo-request	8
time-exceed	11
parameter-problem	12

Source and destination address This feature is somewhat critical from a security point of view, because IP addresses can be reset at any time on a computer. It is not a particularly machine-dependent feature. Source and destination addresses may be specified in different ways:

Type	Example
fully qualified name	www.example.com
IP address	127.0.0.1
network	192.168.10.0/24
	192.168.10.0/255.255.255.0

If source and destination addresses are not specified, the current rule applies to all IP addresses. Normally, private IP addresses are blocked using this method. For a private masked network, such a restriction can also be very useful.

Interface If you specify an interface in a rule, you normally intend this rule to take effect in a specific direction. Rules without specific details of the interface handle all interfaces in the same way. As soon as the direction a packet takes becomes an issue, you should take a closer look at the chain:

	Incoming	Outgoing
PREROUTING	Yes	No
INPUT	Yes	No
FORWARD	Yes	Yes
POSTROUTING	No	Yes
OUTPUT (nat and filter)	No	Yes

Inversion Most criteria allow a negation of their rules. In general, this is done by placing a `!` in front of the corresponding value. For example, `-s !localhost` matches all packets that do not come from the local host.

Subsequent Treatment of Packets (Targets)

After a packet has been successfully identified, a rule must know what it should do with the packet. It is essential that all the following “targets” are written in UPPERCASE.

ACCEPT Pass the packet to the next control point in the diagram.

DROP Drop the packet without generating a return message to the sender.
This is a good target for remote packets that you do not want to accept.

LOG Log the path of the packet.

REJECT Reject the packet as with **DROP**. Normally with **REJECT**, an “ICMP port unreachable” error message is generated, so the sender does not have to wait for a time-out.

RETURN Pass the packet to the default policy.

QUEUE Activate handling by user processes.

MASQUERADE This target is used to masquerade an Internet connection.

REDIRECT Rewrite the packet header. This is mainly used with transparent proxies.

SNAT Changes the source address of the packet.

DNAT Changes the destination address of the packet.

Some of these targets require additional parameters be configured completely. In case of doubt, refer to the man page for `iptables` (`man 8 iptables`).

FreeS/WAN

All data sent over the Internet is routed through unknown computers and networks. If the data is not encrypted, anybody could read it.

Protection against Eavesdropping

There are several approaches to secure the data transfer between companies, subsidiaries, or private persons. At the application layer, this can be achieved with `gpg`, at the transfer layer with `SSL`, and at the network layer with `IPSec`.

The advantage of an implementation at the network layer with `IPSec` is that the encryption is transparent for the applications. This way every access to

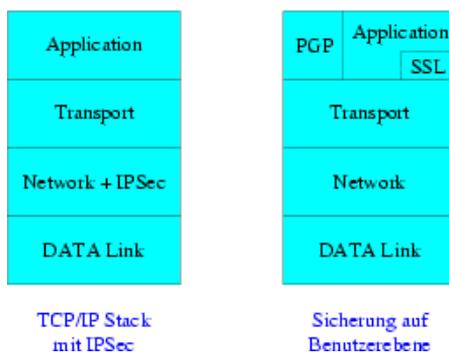


Figure 4.2: Comparing IPSec and SSL

the network can be encrypted as well as the communication between single computers or subnetworks.

It is no problem to connect masqueraded private subnetworks and networks with real Internet IP addresses if the address ranges of the two networks do not overlap. If there are not enough real IP addresses available, masquerade the internal network with private IP addresses as defined in RFC 1918:

```
10.0.0.0      10.255.255.255
172.16.0.0   172.31.255.255
192.168.0.0  192.168.255.255
```

Every member of a VPN must be able to authenticate himself to the other members with a x.509 certificate signed by a CA (Certification Authority). If a certificate is compromised, it must be revoked. It is recommended to organize the VPN like a star with a central host administering the connection between all subnetworks. This host then decides whether a certain member will still be allowed in the VPN.

Security

The encryption of the data packets in FreeS/WAN relies on open algorithms like 3DES. With 3DES, the data is encrypted three times, normally with a 168-bit key. This guarantees a very high level of security, but with enough time and money every existing encryption method can be cracked. This becomes a lot easier if the crackers have access to the private keys or certificates. Make sure neither is available to unauthorized people.

Keys for encryption and decryption can be installed permanently on a host, but it is safer to authenticate with certified keys then negotiate a session key

only valid for a short period. If necessary, a *rekey process* can be started while the connection still exists to replace the old session key.

VPN does not protect against attacks from the Internet. It can prevent unauthorized access to data, but “Denial of Service” attacks and Trojan Horses are still possible. The risk increases if a host communicates with the VPN and the local network. Therefore, the following two recommendations should be taken seriously:

- If the firewall host is also the VPN gateway, it needs to be monitored intensively. No normal users should be able to log in to the firewall host.
- Hosts of remote users, so-called “Road Warriors”, with direct access to the VPN should not be able to connect to the regular LAN. Access to a server in the DMZ (demilitarized zone) should be sufficient.

DNS

The name server BIND Version 8 is used to enable name resolution over the firewall. To learn more about DNS, read the Chapter [DNS — Domain Name Service](#) on page 157 in the appendix of this manual. BIND is configured as a forwarding and caching-only server. All requests will be passed to the forwarders.

Mail

Postfix is a secure, quick, and flexible modular mail transport agent used on the Live CD as a mail relay. A built-in hard disk is absolutely essential for using the mail relay function, because Postfix first has to route the incoming e-mails to a buffer.

HTTP Proxy

To enable highly precise access control to HTTP and HTTPS services, SuSE Live CD for Firewall uses a cascade of different proxies. For access from the inside to the outside and from the outside to the inside, two separate proxy instances are used.

Squid

Squid is an HTTP proxy that offers extensive configuration options. Control over the network clients' access to the web is implemented by means of ACLs (access control lists).

The internal HTTP proxy Squid can be configured to be either transparent or non-transparent. In *non-transparent* mode, the protocols `http`, `https`, and `ftp` are supported. In *transparent* mode, only `http` is supported.

More detailed information about Squid can be found in *Proxy Server: Squid* on page 167.

httpf, tinyproxy

The application `httpf` is responsible for *content filtering*. It is not actually a proxy. Proxy functionality is supplied by the program `tinyproxy`.

It is more specifically a filtering proxy that can prevent downloading and executing program code. This is done by forwarding only known and benign language elements to the web browser. Even the HTTP header entries can be filtered so that, for instance, no information can be posted to the server via the client host's operating system. The configuration can be generated using FAS on the Adminhost.

FTP Proxy

FTP service is split up into two channels. One is the internal connection to the outside. The other is the external connection to the inside or DMZ.

Internal to External

Decide which users of the intranet may access FTP servers in the Internet and via which connections. Normally, the internal FTP proxy remains transparent — the FTP clients from the internal network are automatically forwarded to this proxy, which then sends the requests to the FTP server.

Magic User

Enables the user name, the host, and the destination FTP server's port to provide (e.g., `user@ftp.firma.com:2345`) — automatic execution of the USER command.

Magic Char

Magic char is an option with which the user can specify the destination FTP server and its port himself.

External to Internal

To operate an FTP server, define the settings in this part of the module to enable access from the Internet on the FTP server.

SSH

openssh enables use of a shell on a remote host with an encrypted connection.

chroot, compartment, Kernel Capabilities

To raise the security level on the firewall, the services on the Live CD run in a `chroot` environment. The program `compartment` is also used. Setting capability bits in the kernel additionally increases the security of system applications.

chroot With `chroot`, an application can change its view of the file system irrevocably by defining a new "root" for the file system. As soon as the application has applied itself to this segment of the file system, the segment adopts the role of the entire file system for this application. The rest of the file system no longer exists as far as this application is concerned. Even if the program has somehow crashed, a potential cracker would remain in this `chroot` environment, and so be unable to damage the actual system.

compartment Enables execution of applications and services in `chroot` jails with unprivileged users and groups. It supports scripts run before the program actually starts (e.g., to set up a `chroot` environment). Supports the use of kernel capabilities.

kernel caps Kernel capability bits

Increases security by limiting the capabilities of executable programs. Using `compartment` is another relatively simple option for specifying the capabilities of an application.

The Configuration Disk

The configuration disk contains the complete system and application level gateway configurations. The configuration floppy disk must be formatted

with an ext2 file system and contain the label "SuSE-FWfloppy". Without this label, the configuration floppy is not recognized. The FAS (Firewall Administration System) on the Adminhost creates the file system and the label automatically. The configuration floppy is loaded while the Live CD is booted.

Creating the Configuration Disk

The configuration floppy disk can be created on the Adminhost with FAS. Under 'Configuration', select 'Create floppy disk'. Editing the configuration files with an editor should be left to experts. If necessary, use the editor module of FAS or first create a configuration floppy disk with FAS then modify this according to your needs. *There is no support for this.*

The Configuration Files

The following overview of the configuration files located on the disk is for information purposes only. The files are generated by the Admin desktop:

/etc/hosts See the man page for hosts (man hosts)

/etc/hosts.allow See man 5 hosts_access

/etc/hosts.deny See man 5 hosts_access

/etc/inittab See man inittab

/etc/isdn/ Contains the configuration file for ISDN isdn.conf

/etc/cipe/ Contains the scripts ip-down and ip-up

/etc/live-setup.conf Preparing the hard disk

If you set (*SETUP_DISK=yes*), the hard disk is repartitioned every time the system is booted without asking for confirmation. This means that any existing data is lost. The first partition (/dev/hda1) is used for swap and the second partition (/dev/hda2) for /var.

/etc/ppp/ Directory for ppp configuration

/etc/modules.boot In the file /etc/modules.boot, specify loadable kernel modules to loaded when the system boots. It is not necessary to specify paths relative to /lib/modules/<kernelversion>/. Modules can be given with their names (without a .o file extension), for example, tulip <options> or, with the absolute path, /opt/mymodules/tulip.o.

If necessary, additional options can be passed to the module, such as the IRQ or the IO addresses of the hardware used. Lines beginning with '#' are ignored. If a '-' precedes the module name, an attempt is made to unload the module. This may be necessary if the automatic hardware recognition tries to load the wrong network module, for example.

/etc/ipsec.d/ Contains certificates and configuration files.

/etc/named/ Contains the zone files of the name server.

/etc/named.conf Contains the configuration of bind8.

/etc/named/master/ Contains the master zone files.

/etc/named/slave/ Contains the slave zone files.

/etc/named/root.hint Contains the addresses of the root name server.

/etc/ntp.conf Configuration for the time server daemon xntpd.

/etc/pam.d/ Directory containing the PAM (Pluggable Authentication Module) configuration files.

/etc/permissions.local Sets access permissions for programs and files.

All files are copied from the floppy with the permissions 0600 and root.root. If other permissions are required, they must be specified explicitly here:

```
# Format:
# <file>          <owner>.<group> <permission>
#
/opt/foo/mytool   root.root       755
```

/etc/postfix Configuration directory for Postfix. The major configuration files are

- /etc/postfix/master.cf
- /etc/postfix/main.cf
- /etc/postfix/virtual
- /etc/postfix/transport
- /etc/postfix/access

/etc/proxy-suite/ Configuration directory for FTP proxies.

/etc/rc.config SuSE Linux central configuration file.

/etc/rc.config.d/ This directory contains files used when services are started, for example:

/etc/rc.config.d/i41_hardware.rc.config Configuration of the ISDN card.

/etc/rc.config.d/i41_rc.config Dialing parameters of the ISDN interface.

/etc/rinetd.conf Configuration file for the generic proxy rinetd.

/etc/resolv.conf Configuration file for the resolver library. Includes details of the name server and of the search list.

/etc/route.conf File containing information for generating the static kernel routing table.

/etc/runlevel.firewall Assigns init scripts to the runlevels of the firewall. Init scripts for the corresponding runlevel are started in the order specified. Symbolic links to the scripts are generated. If `network` and `syslog` are specified in `Runlevel 2`:

```
Runlevel:2
network
syslog
```

Links are generated so that, when booting, first `network` then `syslog` is started. The result then appears something like this:

```
/etc/init.d/rc2.d/K84syslog
/etc/init.d/rc2.d/K88network
/etc/init.d/rc2.d/S04network
/etc/init.d/rc2.d/S08syslog
```

Runlevel 2 is the default.

/etc/securetty List of the ttys where the user `root` can log in.

/etc/shadow Contains the encrypted passwords. Normally, these are listed in the form of a `'*'` users except `root`, so only `root` can log in. In FAS, you have the option of not defining a `root` password, in which case access to the firewall would only be possible via SSH and RSA keys.

Caution

No password may be specified for any existing users apart from root. Do not change this file.

Caution

/etc/squid.conf Configuration file for the HTTP proxy Squid.

/etc/ssh/ Contains the configuration files for openssh: `ssh_config` and `sshd_config`.

/etc/syslog.conf Configuration of the syslog daemon. Read the man pages `man 5 syslog.conf`, `man 8 syslogd`, and `man 3 syslog`. The log host and the messages to log are entered in this file. The entry for the log host should appear as follows:

```
 *.* @hostname.domain.tl (or the IP address)
```

/etc/syslog.socks The log daemon `syslog` needs to create a writable socket for all services started in the `chroot` environment. This process is associated with the following files:

```
/var/named/dev/log /var/squid/dev/log
/var/chroot/rinetd/dev/log /var/chroot/ftp-intern/dev/log
/var/chroot/ftp-extern/dev/log
```

/etc/init.d/ This directory contains init scripts specially adapted for the Firewall CD, such as IP packet filter scripts, if any are generated manually or with FAS.

Scripts can also be copied to `/etc/init.d/` on the configuration floppy. Afterwards, the permissions in `/etc/permissions.local` must be specified accordingly, for example, for the init script `/etc/init.d/foobar`:

```
/etc/permissions.local:
#
# Format:
# <file>          <owner>.<group> <permission>
#
/etc/init.d/foobar    root.root          754
```

For the links of the individual runlevels to be created accordingly, add your script to the file `/etc/runlevel.firewall` (here just the line with “foobar” in runlevel 2 and 3 is shown):

/opt the contents of the `/opt` directory on the configuration floppy are copied to the running system in `/opt`. The user can store his files in this directory.

/root/, /root/.ssh/, /root/.ssh/authorized_keys This file contains the RSA public key of the `fwadmin` user on the Adminhost. In the FAS, specify which keys to copy to `/root/.ssh/authorized_keys`. This enables the user `fwadmin` to log in on the firewall as `root`. RSA keys, protected by an additional pass phrase, handle user authentication. This pass phrase is generated during the installation of the Adminhost with the YaST2 module 'Firewall Adminhost'.

Generate the RSA keys at the command line using the command `ssh-keygen(1)`. Read more about it in the man page (`man ssh-keygen`).

It is highly recommended to use the administration desktop FAS on the Adminhost for creating configuration files.

Boot Parameters

Give boot parameters at the `linuxrc` boot prompt as usual (see *Reference*). The following restrictions apply: For security reasons, `init=/bin/bash` and similar parameters cannot be entered at the boot prompt to obtain privileged permissions in a shell. Boot parameters beginning with `init=` are ignored.

IPsec Client on Windows XP and Windows 2000

You can configure the connection manually with the program `ipseccmd.exe` (Windows XP) or `ipsecpol.exe` (Windows 2000), which should be included in your Windows installation. This is a command-line utility, so is very difficult to use. You can also configure the connection with MMC (Microsoft Management Console). However, it is recommended to use `ipsec.exe` for the configuration of the IPsec connection in Windows XP or Windows 2000, since this tool handles the main bulk of the work. This tool can be downloaded from <http://vpn.ebootis.de/package.zip>.

Exporting the Required Certificates

To export the needed certificates, proceed as described here. Open the IPsec module in FAS. Click 'Next' to display your configured VPN connections. Right-click the connection for which to export the certificates. From the drop-down menu, select 'Export certificates'.

Set an appropriate password and remember it. It will be needed when importing the certificates. Save the certificates on an MS-DOS-formatted floppy disk.

Importing the Certificates in Windows

If you use Windows 2000, first install ServicePack2, which enables Windows 2000 to handle "3DES encryption". Otherwise, you will not be able

to connect to Windows 2000. The ServicePack2 is available from <http://www.microsoft.com/windows2000/downloads/servicepacks/sp2/sp2lang.asp>.

For Windows 2000, you also need ipsecpol.exe. Find it in the Resource Kit or at <http://agent.microsoft.com/windows2000/techinfo/reskit/tools/existing/ipsecpol-o.asp>.

Note

Normally, this program is installed to C:/ProgramFiles/ResourceKit. However, since this is a command-line program, copy it to a directory that contains executable files. We recommend copying ipsecpol.exe to C:/WINNT and the respective DLLs to C:/WINNT/System.

Note

Configuring the Required Snap-Ins

Open the MMC by selecting 'Start' → 'Run' and entering mmc. In the MMC, click 'File' → 'Add/Remove Snap-in'. A window displaying active snap-ins opens.

Click 'Add'. A selection window displaying all available snap-ins opens. Select 'Certificates' → 'Add' to start the certificate wizard.

Select 'Computer account'. Click 'Next'. Then select 'Local computer' and click 'Finish'.

Now select the snap-in 'IP Security Policy Management' and click 'Add'. In the certificate wizard, select 'Local computer' then click 'Finish'. Click 'Close' then 'OK'.

Importing the Client Certificate

In the MMC, you will see the two snap-ins that were added. Open the file 'Certificates' by clicking it. Right-click 'Personal Certificates'. In the drop-down menu, select 'All tasks' → 'Import'. This opens the certificate assistant. Click 'Next' then 'Browse'. Under 'File Format', select 'Personal Information Exchange (*.pfx,*.p12)'. Insert the floppy disk containing the exported certificates. Select the floppy drive as source. Select remCert. Click 'Next'.

Now enter the password used in FAS to export the certificates. Click 'Next'. Select the certificate store 'Personal Certificates'. Click 'Next' then 'Finish'.

If the import process was successful, this will be confirmed in a box. Click 'OK'.

Importing a Root Certificate

Right-click 'Trusted Root Certificates'. In the drop-down menu, select 'All Tasks' → 'Import'. This starts the certificate wizard. Click 'Next' then Browse. Under 'File Format', select 'All files (*.*)'. Now select `cacert.pem`. Click 'Next'.

Select the certificate store 'Trusted Root Certification Authorities'. Click 'Next' then 'Finish'. If the import process was successful, it will be confirmed in a box. Click 'OK'.

Making a Note of Important Certificate Data

In the MMC, click 'File' → 'Save'. Save your configuration with the proposed name at the proposed location. In the MMC, open the 'Personal Certificates' folder then the 'Certificates' subfolder.

You will see the certificate for the client. Right-click it. From the drop-down menu, select 'Open'. Select 'Details' and click 'Issuer'. Make a note of the following entries:

```
E=bsupport@suse.de
CN=mainca
OU=bu
O=SuSE
L=Nuernberg
S=Franken
C=DE
```

Close the certificate with 'OK'. Close the MMC with 'File' → 'Close'. When prompted about saving, select 'Yes'.

Configuring the IPsec Connection

Installing the ipsec.exe Tool

Unpack `package.zip` to `C:\ProgramFiles\IPsec\`.

Editing ipsec.conf

Go to the directory C:\ProgramFiles\IPsec. Open the file ipsec.conf with an editor. Adjust the data following the syntax in example 6.

```
conn <name of the connection>
    left=%any
    right=<IP of the Firewall on CD>
    rightsubnet=<IP/netmask of the subnet>
    rightca=<note the previously noted values in the
            reverse order, separated by commas>
    network=auto
    auto=start
    pfs=yes
```

File 6: Syntax of File ipsec.conf

Find an example configuration in 7. Make sure to write the first line left justified and the following lines indented.

```
conn roadwarrior_fwoncd
    left=%any
    right=10.10.254.181
    rightsubnet=192.168.22.0/24
    rightca="C=DE,S=Franken,L=Nuernberg,O=SuSE,OU=bu,
            CN=mainca,E=bsupport@suse.de"
    network=auto
    auto=start
    pfs=yes
```

File 7: Example of an ipsec.conf

Save the edited file.

Creating a Desktop Link and Activating the Connection

If desired, link to the file C:\ProgramFiles\IPsec\IPSEC.exe on your desktop. Now establish the connection to the Internet. Click the created link. A window will open and the IPsec filters will be configured for your current connection. Test the tunnel with ping <client IP behind the tunnel>. The message "IPSec is being negotiated" is displayed once or twice. If the normal ping replies are then displayed, the tunnel is active. In Windows 2000 the second ping call will be successful.

Closing the Connection

To deactivate the IPsec filters and the tunnel, enter `IPSEC.exe -delete`. Create desktop link for this command, if desired.

Implementing the Firewall

On the Adminhost, you created a configuration for the Live CD using FAS or manually created a configuration floppy. In this section, learn how to test this configuration then start the firewall machine.

Requirements for Successful Implementation	138
Booting the Firewall Host	138
Testing the Firewall	138

Requirements for Successful Implementation

First, check to see if your host boots using the configuration set up. Also see if the selected services start and are available for use. Next, check to see if the IP filter of the kernel is working as configured.

Booting the Firewall Host

1. Start the host then open the BIOS setup program. Check the settings for the time and date. It needs to be set to GMT.
2. Configure the boot sequence so the host boots from the CD first — if possible, only from the CD.
3. Assign a BIOS password to prevent, for instance, changes to the boot sequence and to prevent the firewall from being booted from disk.
4. Save the BIOS configuration.
5. Insert the Live CD for Firewall.
6. Insert the configuration floppy.
7. Reboot the host.
8. Note any errors that appear when booting and revise the respective configuration files if necessary. If services do not start (message: "*<Service xyz failed>*"), the corresponding configuration file for the service *<xyz>* contains an error. If this is the case, recreate the configuration disk using the FAS tool on the Adminhost or revise the configuration saved there and rewrite this to the configuration disk.

You should have already defined the configuration of the hard disk using FAS, but if partitions 1 (swap) and 2 (var) of the firewall host do not correspond to the settings in the `/etc/live-setup.conf` file, a dialog screen appears. To reconfigure the hard disk, answer 'Yes'.

Testing the Firewall

Before employing the firewall for daily use, testing should be done to ensure the packet filter is configured correctly and the (configured) proxies all start

and process requests properly. Adminhost tools are available on the firewall for these purposes (see 2 on page 15), such as `nmap`, `nessus`, `xlogmaster`, `logsurfer`, and `http` clients. Detailed documentation for these programs is available in the `/usr/share/doc/packages/` directory on the SuSE Adminhost for Firewall. Man pages are also available for each program.

Only when every single test has been successfully completed can you start the firewall. Document all tests conducted.

First, connect a laptop or other computer to the firewall host to simulate an external network. Then, close the connection to the internal network and establish a connection to the Internet. If possible, test your firewall from the outside. Check your setup.

Internal Testing

Tests to conduct:

- Are all services available?
- Test if the permitted services are working from the internal clients. Can you access `https`, send e-mails, and transfer data via `FTP`?
- Do the deny rules work?
- Test your packet filter. Try using a port scanner like `nmap`. Follow the log messages of your firewall on the log host or on your firewall itself. Let a packet sniffer run simultaneously to detect restricted packets or to see if response packets are not being sent.
- Are the log files being written to the log host? If the Adminhost is also the log host and the feature has been activated in FAS, use FAS for analyzing the logs conveniently.

Fixing errors:

Determine the source of the problem. Search for the log files according to process name, for example, `postfix` or `named`:

```
earth:~ # grep postfix /var/log/messages
```

or

```
earth:~ # grep named /var/log/messages
```

Alternatively, use the search dialog of the FAS Log File Analysis module. Many programs can be switched to “verbose mode”. In this way, obtain detailed information, which can, however, be quite extensive.

External Testing

Test externally to see if the available services are working. For instance, if you can send e-mails to the internal network. You should be able to see in the postfix messages in `/var/log/mail` on the firewall host whether the e-mails were accepted and could be delivered to the internal mail server. Check to see if the packet filter is working. This can be verified by a port scanner. At the same time, find the kernel packet filter messages in `/var/log/messages` on the firewall machine as well as in the log file analyzer in FAS. Try to set up connections to explicitly restricted ports and attempt to find the corresponding log entries and match them to their corresponding events. For example, try `telnet` to port 79 of the firewall. This should be logged on the firewall and the log host.

If you are using a log host, check to see if the log messages are being transmitted in their entirety. If the FAS Log File Analysis module is configured accordingly, it can be used for this purpose.

Going Online

Only after you have completed all these tests, connect the firewall host to the Internet and your intranet and begin productive operation.

Note

Constantly monitor your log files. This is the only way to ensure a timely response to attacks or failures. If unusual events occur, react immediately. Use the FAS Log File Analysis module to check and track the log files.

Note

Help

In this chapter, find information about creating a setup concept for a firewall solution in your network using the SuSE Firewall on CD. Also find information about using the services of SuSE Linux AG to have a plan drawn up tailored to your specific needs.

Troubleshooting	142
Detecting Attacks	143
Recommended Reading	148

Troubleshooting

Find help here if the Adminhost cannot be installed or if the Live CD is not booting.

Problems Installing the Adminhost

If you have trouble installing the Adminhost, SuSE Linux Installation Support is available free of charge. First, check the support database. Here, find a realm of information about a variety of installation problems. A keyword search helps find relevant information. The support database is available at <http://sdb.suse.de>.

Problems Booting the Live CD

To simplify your troubleshooting, observe and note the following:

- What happens when you boot?
- Are services not started that have already been configured (skipped or failed messages)?
- Do error messages appear on the console?
- Are there kernel error messages on virtual consoles 9 or 12? To see these messages, change to this console using `(Alt) + (F9)` or `(Alt) + (F12)`.

Problems Integrating the Network

Observe and note the nature of the problems. For example, on the firewall, test whether the network interfaces are configured correctly (log in to the firewall as user `root`) by entering `ifconfig`.

This command shows which interfaces are configured and which IP addresses, network masks, and other details are set up. If necessary, correct the configuration using FAS and create a new configuration floppy. The host must be restarted to commit the new configuration.

- Which services function from the client side (intranet)? Should they be functioning? The log files can help determine whether it is a problem involving unauthorized access. Try reproducing the errors.

- Is external access to available resources not functioning? Which services are affected? Should the resources really be accessible?

Test, using `ps`, whether the process is available — if it can be accessed. If services are not accessible, check your log files. Look for messages about why a particular service was not started or whether an unauthorized party has been attempting to make use of the service. Test from several clients whether the firewall host is accessible and whether the proxies are responding.

Detecting Attacks

Intrusion Detection and Event Display

A properly configured Linux/UNIX system can, in and of itself, be considered quite secure. Internal system hazards associated with a complex system such as Linux or UNIX are more easily recognized than on other operating systems, because UNIX has been used and developed for over thirty years. UNIX also forms the basis of the Internet. Nevertheless, configuration errors can occur and security holes can appear. There will always be security flaws. Security experts and crackers are in perpetual competition to be one step ahead of the other. What qualifies today as secure may be vulnerable tomorrow.

Signs of Intrusion

Any abnormal behavior on your firewall system can be considered a sign that your system is compromised:

- increased processor load
- unusually heavy network traffic
- unusual processes
- processes started by nonexistent users

Recognizing an Intrusion

First, understand which actions are defined as intrusions. Unfortunately, it is normal these days that a host connected to the Internet will be scanned for open and vulnerable ports. It is just as common that ports recognized as vulnerable (such as POP3, qpopper, rpc-mountd, smtp, or sendmail) are attacked. Most of these attacks are carried out by “script kiddies”. Pre-fabricated “exploits” published on relevant web sites (e.g., <http://www.rootshell.com>) are used. These web pages also exist as valuable information sources for network and system administrators. They can usually be identified by the fact that the hack is only carried out once and is not repeated if the action fails on the first attempt. The first measure to take in case of such an event is to check if a break-in really did happen. Furthermore, raise the log level, for example by logging accepted packets (expert configuration of IP filter and NAT module), and refine the evaluation (e.g., a selective search aimed at an intruding IP address in the log files or any unusual port numbers).

Determine for yourself what is a dangerous attack on your system. At least determine how to react to such an event. The following literature is recommended if you suspect an intrusion: “Steps for Recovering from a Unix Root Compromise” (http://www.cert.org/tech_tips/root_compromise.html) and RFC 2196 Site Security Handbook. Both publications describe procedures that should follow any successful break-in. The documents provide a formal starting point as to how a company, agency, or educational institution can react. The procedure discussed there necessitates a certain amount of memory to take “snapshots” of the system and requires colleagues to carry out the analysis and to examine the security problem. Situations are described that may necessitate taking punitive action.

Responding to an Intrusion

If you suspect a successful attack:

- It is important to stay calm. Hasty actions can destroy important information (e.g., processes initiated by the cracker that contain information about what the cracker is doing or how he is attacking).
- The course of action, as well as who to inform, should be outlined in the security policy. The communication should not take place by e-mail, but by telephone or fax.
- Physically cut the network connection to the firewall. It is not a good idea to shut down the computer. Important information could be lost, such as programs started by the cracker.

- List all running processes with `ps` and search for processes that do not typically occur in normal firewall operation.
- Draw up a process table when setting up the firewall that can later serve as a basis for comparison.
- Check the running processes for links to unusual TCP or UDP ports.
- See if the packet filter rules were changed.
- Compare all configuration files with the original configuration. This is very easy with the SuSE Firewall on CD, because the configuration is already saved on the Adminhost. In doing so, however, it is essential not to restart the firewall host, because any changes to the filter rules could be lost.
- Also back up all log files. The log files could be legally admissible evidence. Document all steps you taken. If you save the log files locally to the hard disk, make an exact copy of the hard disk for documentation purposes.
- Save your data to CD or to another medium (tape drive, ZIP drive).
- Analyze the log files: Who tried to access what services or ports when and from where (IP address, domain name, possibly even user name)? Was an attempt made to uncover passwords (multiple failed login attempts with the same user name)?
- Make sure you are using a uniform and exact time source. It is important that the hardware clock of the firewall host and the log host are as closely in sync as possible. Use a common time source for all your servers whenever possible. Only in this way can you keep track of events accurately.
- Which parts of the security policy were violated? This is especially important in regard to internal intrusions.
- You may want to deactivate the user account in your network that carried out the attack, if it was an internal attack. The relevant procedure pertaining to internal violations should also be regulated (security or company policy).

External Attacks

Inform the system administrator responsible for the address block (via postmaster or the domain's abuse address). The report of an incident or attack should contain enough information to ensure that the other party can investigate the problem. However, consider that your contact person could be the one who has carried out the attack. Here is a list of possible information to provide. Decide which of the following pieces of information to give:

- Your e-mail address
- Telephone number
- Your IP address, host name, domain name
- The IP addresses, host names, and domain names affected by the hacking incident
- The date and time of the intrusion, preferably with the time zone
- A description of the attack
- Explain how the attack was recognized
- Excerpts of the log files relating to the attack
- A description of the log file format
- Details of advisories and security information that describes the nature and severity of the attack
- What you want the contact person to do: Close an account, confirm the occurrence of an attack, issue a report for information purposes only, request for further observation

Once you have gone through all the data security and documentation procedures, set up your firewall again. Raise the log level of each application if possible. It is likely the cracker will try to infiltrate your system again. This will be the opportunity to catch the intruder red-handed.

Examples:

Log in to the console.

Examine the log files for messages of the IP filter.

Search for certain unusual IP addresses (frequently occurring rejection of packets that correspond to IP addresses on one or more port numbers). Find out exactly what happened. Using FAS, examine the log files according to definable criteria.

Sometimes, it is unclear what to look for until you find it. It is also possible that a system administrator from another network is reporting a complaint, via postmaster or abuse mail address, that attacks have been occurring from your network to remote hosts. Take such complaints seriously. Request that logs of all intrusions are sent so you have an idea of the date and time as well as the method of attack on the external system. Attempt to verify the circumstances surrounding the incident. A cracker, intruder, or attacker may have already overstepped the security boundaries and misused your network. Here as well, the reputation of your business is at stake.

Once you have backed everything up and documented it, review your firewall configuration. After fixing possible errors or after shutting down services at risk, restart your firewall. This shows the clear advantage of the SuSE Firewall on CD: the original status of the firewall can be restored simply by rebooting the firewall host, making a complicated reinstallation of the operating system and restoring backups unnecessary.

Advantage of the Live File System of the SuSE Firewall on CD

One of the greatest advantages of the SuSE Firewall on CD is that its initial state can be restored simply by booting it. Keep in mind that any possible configuration errors may reappear.

If the firewall has been breached, the method of intrusion should be investigated to correct configuration errors. If security holes are found in applications, SuSE Linux AG provides updates for the affected applications — in the case of the Live CD, a new CD.

Find more information at <http://www.suse.de/de/security/>, <http://www.cert.org>, and <http://www.first.org>.

Recommended Reading

- D. Brent Chapman & Elizabeth D. Zwicky: *Building Internet Firewalls*, 2nd edition, O'Reilly 2000.
- *Maximum Linux Security*, SAMS 1999.
- Robert L. Ziegler: *Linux Firewalls*, New Riders 1999.

Support, Maintenance, and Patch Management

Maintenance

With SuSE Maintenance, always have the most up-to-date patches and security fixes for your SuSE Linux Business product available. For your system to remain up-to-date, you need to check regularly for new patches on the SuSE Maintenance Web.

Your purchase of SuSE Firewall on CD includes protected access to the SuSE Maintenance Web for twelve months after the date of registration.

Accessing the SuSE Maintenance Web

Register your product online at <http://support.suse.de/en/register/>. You will then receive an e-mail with instructions on how to proceed. The “SuSE Maintenance Web” can be reached at <http://support.suse.de/psdb/>.

Getting Patches

Patches can be downloaded from the Maintenance Web both for the Admin CD and for the Live CD.

Patches for the Admin CD

There are two possibilities here:

Via the SuSE Maintenance Web

Log in to the SuSE Maintenance Web and download patches individually (see the above URL). Under 'Products', select the SuSE Firewall on CD— Adminhost or, if you have installed the VPN module, Adminhost with VPN, then choose the needed patches. Detailed documentation is available for each patch.

Using YOU (YaST Online Update)

The SuSE Tool YOU establishes a connection to the SuSE Maintenance Web and checks to see which patches are relevant to your system. These packages are then installed automatically by YOU on your Adminhost. The log is then displayed, if desired.

Start YOU from KDE with the YaST2 control center or from a console window (like xterm) by logging in as the user root with `sux -` then entering `yast2`.

Note

YOU uses the software package `wget` to download patches. If this package is not yet installed on your Adminhost, install it first.

Note

Check to see if package `wget` is installed with `rpm -q wget`. If not, Use YaST2 ('YaST2' → 'Install software' → 'Search' then enter `wget`) to install the package.

If you make a direct connection to the outside, no further settings are necessary. If the connection is established via a proxy, the file `/etc/wgetrc` must be adjusted. Open the file `/etc/wgetrc` with an editor of your choice. Enter the following line:

```
http_proxy = http://<your_proxy>:<port>/
```

`<your_proxy>` is the name of your proxy server and `<port>` is the port number to which the proxy is listening. Specify the HTTP proxy to use, not the FTP proxy. Patches are always downloaded via HTTP, not via FTP.

Patches for the Live CD

The patches for the Live CD are made available in the form of an RPM package. This RPM package contains an ISO file (the new Live CD). As with the

Admin CD, download the patches via the SuSE Maintenance Web and burn the ISO file to a CD. Find instructions for the patch on the SuSE Maintenance Web. This procedure has the great advantage that your Live CD is absolutely up-to-date.

As soon as a new ISO file appears for the Live CD, you will be sent a new Live CD, regardless of whether you have already downloaded the file. If several ISO files appear within a short period of time, for example, because several different security holes have been discovered at short intervals, you will be sent the most recent Live CD to appear. Because it takes a number of days to produce and send the CD, it would not make sense to send several no longer up-to-date CDs.

Burning ISO Images with `cdrecord`

An ISO image is an “image” of a CD. This image can be mounted or burnt to a CD. For reasons of security, you can only burn CDs in Linux as the user `root`. Change to `root` with `su`. After entering the root password, check to which bus the CD burner is connected with `cdrecord -scanbus`. This generates output containing something like:

```
...
scsibus0:
  0,0,0  0) *
  0,1,0  1) *
  0,2,0  2) *
  0,3,0  3) *
  0,4,0  4) 'YAMAHA  ' 'CRW8824S          ' '1.00' Removable CD-ROM
  0,5,0  5) *
  0,6,0  6) 'TOSHIBA ' 'DVD-ROM SD-M1201' '1011' Removable CD-ROM
  0,7,0  7) *
```

In this case, the CD burner was made by Yamaha. The numbers 0,4,0 specify the target device. Your numbers may be different, depending on the configuration. Make note of those numbers. To burn the ISO image to CD, enter the following as the user `root`:

```
root@earth:~ > cdrecord -v dev=X,Y,Z speed=4 -eject
```

In our example, replace the place holders X, Y, and Z with the correct values, noted earlier. The options here have the following meanings:

`-v` provides detailed messages

`dev` specifies the SCSI device of the burner as a number

`speed` sets the speed of the burning process

`-eject` ejects the CD after burning is completed

Find more options in the man page for `cdrecord` (`man cdrecord`).

Support and Services

Support Conditions

Scope of Installation Support

The installation support can assist with installing your SuSE Linux system. This also applies to the central system components that allow for fundamental operation. It includes:

- Installation of the SuSE Linux base system on a host from the “Admin CD for Firewall”.
- The configuration of the basic hardware of this host with the graphical installation tool YaST2— the central components of the PC, excluding peripheral devices, but including the configuration of an ethernet card.

Topics not mentioned here are not covered by installation support.

Duration of Installation Support

The installation support for the Admin CD for Firewall lasts for a period of 30 days following the registration date.

Reaching the SuSE Support Team

Reach our support team by e-mail, fax, or mail:

- **E-mail:**
Address: fw-support@suse.de
Processing: weekdays
- **Fax:**
Fax number: +49-09 11-74 05 34 89
Processing: weekdays

- **Mail:**
 - Address: SuSE Linux AG
— Support —
Deutschhernnstr. 15-19
D-90429 Nürnberg
 - Processing: weekdays

Commercial Support

Even if an operating system comes with all the necessary facilities, it will only be a viable alternative for use in the corporate environment in combination with professional and qualified support services. SuSE guarantees this kind of service for Linux. All information about this can be found at the central support portal for SuSE Linux:

<http://www.suse.de/en/business/services/support>

Individual Projects and Consulting

SuSE Linux AG offers competent consulting and solutions for your individual needs. We have a great deal of experience in the deployment of Linux servers. You can make use of the know-how of our experts to successfully realize your projects.

Our strength lies in our versatility. Databases, security issues, Internet connections, or company-wide networks: with the right software, Linux is a powerful platform for your applications. Our services range from the conception, implementation, and configuration of server systems to a complete infrastructure consultation.

You may want, for instance, to implement your Internet presence on a SuSE Linux system and are therefore looking for web server, e-mail, and secure server solutions. Our consultants can help you design and implement the right solution.

Are you managing a complex heterogeneous network in which you would like to integrate Linux? We offer consultation and support for the design and implementation of complex server solutions.

Do you have special requirements or needs that are not fulfilled by standard software? We can assist you further in individualized system development.

SuSE Linux AG
Deutschhernnstr. 15-19

D-90429 Nürnberg
Tel: +49-911-740-53-0
Fax: +49-911-740-53-479
E-mail: suse@suse.de

Represented by our Regional Service Centers in Germany and outside, as well as our Support and Development Center in Nuremberg, on-site support is provided by the company SuSE Linux AG:

- Rollout and Implementation Services
- Infrastructure Consultation
- Intranet Server Solutions
- Internet Server Solutions
- Development of Client-Specific Requirements
- Complete Solutions
- E-Commerce

SuSE Training Program

SuSE provides training courses and workshops for Linux. Our comprehensive program ranges from user courses through administration training to courses for developers and preparation courses for LPI certification. Only SuSE trainers or those certified by SuSE Training (SCLT), who have met our high standards, may teach our courses, ensuring the highest levels of expertise. To guarantee a high rate of knowledge transfer, we place great value on up-to-date and practically-oriented material, tailored to classroom needs.

With the SuSE Training program, develop and reinforce the knowledge a successful company in the IT field requires. Schedule a course with one of our certified training partners close to you. Our trainers can also come to you to train your staff individually in courses tailored to your needs. A continually updated overview of all course dates, together with a detailed description of course contents, can be found at:

<http://www.suse.de/en/business/services/support/training>

Feedback

We always appreciate your tips, hints, and problem descriptions. We will help you if your problem is a straightforward one or if we already have the solution at hand.

SuSE makes every effort to construct a Linux system that meets the wishes of our customers as closely as possible. We therefore appreciate any criticism of our CD or of this book as well as suggestions for future projects. We think this the best way to correct errors and to maintain the high quality standards of SuSE Linux.

Send feedback any time through our web site: www.suse.de/feedback.

Additional Services

We would also like to draw your attention to our services available at all times:

- **SuSE WWW Server**

<http://www.suse.com>

Up-to-date information, catalogs, ordering service, support form, and support database

- **SuSE Mailing Lists** (information and discussions via e-mail):

- ▷ suse-announce-e@suse.com — announcements concerning SuSE GmbH
- ▷ suse-linux-e@suse.com — Discussions all about SuSE Linux
- ▷ proxy-suite@suse.com — Discussion concerning the SuSE proxy suite
- ▷ suse-axp@suse.com — SuSE Linux on alpha processors
- ▷ suse-ham-e@suse.com — SuSE Linux and amateur radio
- ▷ suse-ibm-db2@suse.com — SuSE Linux and IBM DB2
- ▷ suse-informix@suse.com — Information and discussion concerning Informix in SuSE Linux
- ▷ suse-motif@suse.com — SuSE Linux on Motif
- ▷ suse-oracle@suse.com — Information and discussion concerning Oracle in SuSE Linux
- ▷ suse-ppc@suse.com — SuSE Linux on Power PC processors

- ▷ suse-security@suse.com — Discussion of security issues in SuSE Linux
- ▷ suse-security-announce@suse.com — Announcement of security-related errors and updates
- ▷ suse-sparc@suse.com — SuSE Linux on Sparc processors

To subscribe to a list, send an e-mail message to `<LISTNAME>-subscribe@suse.com`.

An automatic response will be sent back that you will need to confirm. For `<LISTNAME>`, substitute the name of the mailing list to which to subscribe, for example, suse-announce-subscribe@suse.com to receive regular announcements.

The procedure is similar for unsubscribing from a list: `<LISTNAME>-unsubscribe@suse.com`. Send the unsubscribe mail with e-mail address that is subscribed.

■ SuSE FTP Server

<ftp://ftp.suse.com>

Current information, updates, and bug fixes. Log in to the system as user `ftp`.

DNS — Domain Name Service

DNS (Domain Name Service) is needed to resolve domain and host names into IP addresses. This chapter describes how to configure the name server BIND9. It includes information about the configuration files `named.conf`.

Starting the Name Server BIND	158
The Configuration File <code>/etc/named.conf</code>	159
For More Information	166

Starting the Name Server BIND

The name server BIND is already preconfigured in SuSE Linux, so you can easily start it right after installing the distribution.

If you already have a functioning Internet connection and have entered 127.0.0.1 as name server for the local host in `/etc/resolv.conf`, you should normally already have a working name resolution without having to know the DNS of the provider. BIND carries out the name resolution via the root name server, a notably slower process. Normally, the DNS of the provider should be entered with its IP address in the configuration file `/etc/named.conf` under forwarders to ensure effective and secure name resolution. If this works so far, the name server will run as a pure “caching-only” name server. Only when you configure its own zones will it become a proper DNS. A simple example of this can be found under `/usr/share/doc/packages/bind8/sample-config`. However, do not set up any official domains until assigned one by the responsible institution. Even if you have your own domain and it is managed by the provider, you are better off not to use it, as BIND would otherwise not forward any more requests for this domain. The provider’s web server, for example, would not be accessible for this domain.

To start the name server, enter `rcnamed start` at the command line as root. If “done” appears to the right in green, `named`, as the name server process is called, has been started successfully. Immediately test the functionality of the name server on the local system with the `nslookup` program. The local host should appear as the default server with the address 127.0.0.1. If this is not the case, the wrong name server has probably been entered in `/etc/resolv.conf` or this file does not exist. For the first test, enter `nslookup “localhost”` or `“127.0.0.1”` at the prompt, which should always work. If you receive an error message instead, such as “No response from server”, check to see if `named` is actually running using the command `rcnamed status`. If the name server is not starting or is exhibiting faulty behavior, find the possible causes of this logged in `/var/log/messages`.

If you have a dial-up connection, be sure that BIND8, once it starts, will re-view the root name server. If it does not manage this because an Internet connection has not been made, this can cause the DNS requests not to be resolved other than for locally-defined zones. BIND9 behaves differently, but requires quite a bit more resources than BIND8.

To implement the name server of the provider or one already running on your network as “forwarder”, enter one or more of these in the options section under forwarders. See File 8.

```
options {
    directory "/var/lib/named";
    forwarders { 10.11.12.13; 10.11.12.14; };
    listen-on { 127.0.0.1; 192.168.0.99; };
    allow-query { 127/8; 192.168.0/24; };
    notify no;
};
```

File 8: Forwarding Options in named.conf

Adjust the IP addresses to your personal environment. After options follows the zone, “localhost”, “0.0.127.in-addr.arpa”, and “.” entries. At least entries from “type hint” should exist. Their corresponding files never have to be modified, as they function in their present state. Also, be sure that a “;” follows each entry and that the curly braces are properly set.

If you have made changes to the configuration file `/etc/named.conf` or to the zone files, have BIND reread these files by entering `rndc reload`. Otherwise, completely restart the name server with `rndc restart`. To stop the name server, enter `rndc stop`.

The Configuration File `/etc/named.conf`

Make all the settings for the name server BIND8 and BIND9 in the `/etc/named.conf` file. The zone data, consisting of the host names, IP addresses, and similar, for the domains to administer are stored in separate files in the `/var/lib/named` directory.

The `/etc/named.conf` is roughly divided into two areas. One is the options section for general settings and the other consists of zone entries for the individual domains. Additional sections for logging and acl type entries can be added. Comment lines begin with a ‘#’ sign or ‘//’. A minimalistic `/etc/named.conf` looks like File 9.

```
options {
    directory "/var/lib/named";
    forwarders 10.0.0.1; ;
    notify no;
};

zone "localhost" in {
    type master;
    file "localhost.zone";
```

```

};

zone "0.0.127.in-addr.arpa" in {
    type master;
    file "127.0.0.zone";
};

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

```

File 9: A Basic /etc/named.conf

This example works for both BIND8 and BIND9, because no special options are used that are only understood by one version or the other. BIND9 accepts all BIND8 configurations and makes note of options not implemented at start-up. Special BIND9 options are, however, not supported by BIND8.

Important Configuration Options

directory "/var/lib/named"; specifies the directory where BIND can find the files containing the zone data.

forwarders 10.0.0.1; ; is used to specify the name servers (mostly of the provider) to which DNS requests, which cannot be resolved directly, are forwarded.

forward first; causes DNS requests to be forwarded before an attempt is made to resolve them via the root name servers. Instead of `forward first`, `forward only` can be written to have all requests forwarded and none sent to the root name servers. This makes sense for firewall configurations.

listen-on port 53 127.0.0.1; 192.168.0.1; ; tells BIND to which network interface and port to listen. The port 53 specification can be left out, since 53 is the default port. If this entry is completely omitted, BIND accepts requests on all interfaces.

query-source address * port 53; This entry is necessary if a firewall is blocking external DNS requests. This tells BIND to post requests externally from port 53 and not from any of the ports greater than 1024.

allow-query 127.0.0.1; 192.168.1/24; ; defines the networks from which clients can post DNS requests. The /24 at the end is an abbreviated expression for the netmask, in this case 255.255.255.0.

allow-transfer ! *; ; controls which hosts can request zone transfers. This example cuts them off completely due to the ! *. Without this entry, zone transfers can be requested anywhere without restrictions.

statistics-interval 0; In the absence of this entry, BIND8 generates several lines of statistical information in `/var/log/messages`. Specifying 0 suppresses these completely. Otherwise the time in minutes can be given here.

cleaning-interval 720; This option defines at which time intervals BIND8 clears its cache. This triggers an entry in `/var/log/messages` each time it occurs. The time specification is in minutes. The default is 60 minutes.

interface-interval 0; BIND8 regularly searches the network interfaces for new or no longer existing interfaces. If this value is set to 0, this will not be carried out and BIND8 will only listen at the interfaces detected at start-up. Otherwise, the interval can be defined in minutes. The default is 60 minutes.

notify no; no prevents other name servers from being informed when changes are made to the zone data or when the name server is restarted.

The Configuration Section “Logging”

What, how, and where archiving takes place can be extensively configured in BIND8. Normally, the default settings should be sufficient. File 10 represents the simplest form of such an entry and will completely suppress any logging:

```
logging {  
    category default { null; };  
};
```

File 10: Entry to Suppress Logging

Zone Entry Structure

```
zone "my-domain.de" in
{
    type master;
    file "my-domain.zone";
    notify no;
};
```

File 11: Zone Entry for my-domain.de

After zone, the name of the domain to administer is specified, my-domain.de, followed by in and a block of relevant options enclosed in curly braces, as shown in File 11. To define a “slave zone”, the type is simply switched to slave and a name server is specified that administers this zone as master (but can also be a “slave”), as shown in File 12.

```
zone "other-domain.de" in {
    type slave;
    file "slave/other-domain.zone";
    masters { 10.0.0.1; };
};
```

File 12: Zone Entry for other-domain.de

The options:

type master; master indicates that this zone is administered on this name server. This assumes that your zone file has been properly created.

type slave; This zone is transferred from another name server. Must be used together with masters.

type hint; The zone . of the type hint is used for specification of the root name servers. This zone definition can be left alone.

file “my-domain.zone” or file “slave/other-domain.zone”; This entry specifies the file where zone data for the domain is located. This file is not required by slaves, because its contents is read by another name server. To differentiate master and slave files, the directory slave is specified for the slave files.

masters { 10.0.0.1; }; This entry is only needed for slave zones. It specifies from which name server the zone file should be transferred.

allow-update { ! *; }; This options controls external write access, which would allow clients to make a DNS entry — something which is normally not desirable for security reasons. Without this entry, zone updates are not allowed at all. Note that with the above sample entry, the same would be achieved because ! * effectively bars any clients from such access.

Structure of Zone Files

Two types of zone files are needed: one serves to assign IP addresses to host names and the other does the reverse — supplies a host name for an IP address.

'.' has an important meaning in the zone files here. If host names are given without ending with a '.', the zone will be appended. Thus, complete host names specified with a complete domain must end with a '.' so the domain is not added to it again. A missing point or one in the wrong place is probably the most frequent cause of name server configuration errors.

The first case to consider is the zone file `world.zone`, responsible for the domain `world.cosmos`, as in File 13.

```

1. $TTL 2D
2. world.cosmos.  IN SOA      gateway  root.world.cosmos. (
3.                2001040901 ; serial
4.                1D        ; refresh
5.                2H        ; retry
6.                1W        ; expiry
7.                2D )      ; minimum
8.
9.                IN NS      gateway
10.               IN MX      10 sun
11.
12. gateway       IN A       192.168.0.1
13.               IN A       192.168.1.1
14. sun           IN A       192.168.0.2
15. moon          IN A       192.168.0.3
16. earth         IN A       192.168.1.2
17. mars          IN A       192.168.1.3

```

File 13: The File `/var/lib/named/world.zone`

Line 1: \$TTL defines the standard TTL that applies for all the entries in this file, here 2 days. TTL means “time to live”.

Line 2: The SOA control record begins here:

- The name of the domain to administer is world.cosmos in the first position. This ends with a `.` , because otherwise the zone would be appended a second time. Alternatively, a `@` can be entered here. Then, the zone would be extracted from the corresponding entry in `/etc/named.conf`.
- After IN SOA is the name of the name server in charge as master for this zone. The name is extended from gateway to gateway.world.cosmos, because it does not end with a `.` .
- Afterwards, an e-mail address of the person in charge of this name server will follow. Since the `@` sign already has a special significance, `.` is to be entered here instead, for root@world.cosmos, consequently root.world.cosmos..The `.` sign at the end cannot be neglected, otherwise, the zone will still be added here.
- A `(` follows at the end here, including the following lines up until `)` into the SOA record.

Line 3: The serial number is an arbitrary number that is increased each time this file is changed. It is needed to inform the secondary name servers (slave servers) of changes. For this, a ten-digit number of the date and run number, written as YYYYMMDDNN, has become the customary format.

Line 4: The refresh rate specifies the time interval at which the secondary name servers verify the zone serial number. In this case, 1 day.

Line 5: The retry rate specifies the time interval at which a secondary name server, in case of error, attempts to contact the primary server again. Here, 2 hours.

Line 6: The expiration time specifies the time frame after which a secondary name server discards the cached data if it has not regained contact to the primary server. Here, it is a week.

Line 7: The minimum time to live states how long the results of the DNS requests from other servers can be cached before they become invalid and have to be requested again.

Line 9: The IN NS specifies the name server responsible for this domain. The same is true here that gateway is extended to gateway.world.cosmos because it does not end with a `.`. There can be several lines like this, one for the primary and one for each secondary name server. If notify is not set to no in `/etc/named.conf`, all the name servers listed here will be informed of the changes made to the zone data.

Line 10: The MX record specifies the mail server that accepts, processes, and forwards e-mails for the domain world.cosmos. In this example, this is the host sun.world.cosmos. The number in front of the host name is the preference value. If there are multiple MX entries, the mail server with the smallest value is taken first and, if mail delivery to this server fails, an attempt will be made with the next higher value.

Line 12–17: These are now the actual address records where one or more IP addresses are assigned to the host names. The names are listed here without a `.` , because they are entered without a domain added and can all be appended with world.cosmos. Two IP addresses are assigned to the host gateway, because it has two network cards.

The pseudodomain in-addr.arpa is used to assist the reverse lookup of IP addresses into host names. This will be appended, for this purpose, to the network components described here in reverse order. 192.168.1 is thus translated into 1.168.192.in-addr.arpa. See File 14.

```

1. $TTL 2D
2. 1.168.192.in-addr.arpa. IN SOA  gateway.world.cosmos.
                               root.world.cosmos. (
3.                               2001040901      ; serial
4.                               1D              ; refresh
5.                               2H              ; retry
6.                               1W              ; expiry
7.                               2D )           ; minimum
8.
9.                               IN NS         gateway.world.cosmos.
10.
11. 1.                          IN PTR       gateway.world.cosmos.
12. 2.                          IN PTR       earth.world.cosmos.
13. 3.                          IN PTR       mars.world.cosmos.
```

File 14: Reverse Lookup

Line 1: \$TTL defines the standard TTL that applies to all entries here.

Line 2: 'Reverse lookup' should be activated with this file for the network 192.168.1.0. Since the zone is called '1.168.192.in-addr.arpa' here, it is, of course, undesirable to add this to the host name. Therefore, these are all entered complete with domain and ending with '.'. The rest corresponds to the previous example described for world.cosmos.

Line 3–7: See the previous example for world.cosmos.

Line 9: This line also specifies the name server responsible for this zone. This time, however, the name is entered completely with domain and ending with '.'.

Line 11–13: These are the pointer records which are linked to an IP address at the respective host name. Only the last part of the IP address is entered at the beginning of the line missing the last '.'. Now, if the zone is appended to this and the .in-addr.arpa is neglected, the entire IP address will be backwards.

In this form, the zone files are usable both for BIND8 and BIND9. Zone transfers between different versions should not normally be an issue.

For More Information

- Documentation on package bind8: <file:///usr/share/doc/packages/bind8/html/index.html>.
- A sample configuration can be found at:
`/usr/share/doc/packages/bind8/sample-config`
- the man page for named (man 8 named) in which the relevant RFCs are named and the man page for named.conf (man named.conf)

Proxy Server: Squid

The following chapter describes how caching web sites assisted by a proxy server works and what the advantages of using Squid are. The most popular proxy cache for Linux and UNIX platforms is Squid. We will discuss its configuration, the specifications required to get it running, how to configure the system to do transparent proxying, how to gather statistics about the cache's use with the help of programs like Calamaris and cachemgr, and how to filter web contents with squidgrd.

What is a Proxy Cache?	168
Some Facts About Cache Proxying	168
System Requirements	170
Starting Squid	172
The Configuration File <code>/etc/squid.conf</code>	173
Transparent Proxy Configuration	178
Squid and Other Programs	179
More Information on Squid	182

What is a Proxy Cache?

Squid acts as a proxy cache. It behaves like an agent that receives requests from clients, in this case web browsers, and passes them to the specified server provider. When the requested objects arrive at the agent, it stores a copy in a disk cache.

Benefits arise when different clients request the same objects: these will be served directly from the disk cache, much faster than obtaining them from the Internet, and, at the same time, saving overall bandwidth for the system.

Tip

Squid covers a wide range of features, including intercommunicating hierarchies of proxy servers to divide the load, defining strict access control lists to all clients willing to access the proxy, and, with the help of other applications, allowing or denying access to specific web pages. It also can obtain statistics about the most visited web sites, user usage of the Internet, and others.

Tip

Squid is not a generic proxy. It proxies normally only between HTTP connections. It does also support the protocols FTP, Gopher, SSL, and WAIS, but it does not support other Internet protocols such as Real Audio, news, or videoconferencing. Because Squid only supports the UDP protocol to provide communication between different caches, many other multimedia programs will not be supported.

Some Facts About Cache Proxying

Squid and Security

It is also possible to use Squid together with a firewall to secure internal networks from the outside using a proxy cache. The firewall denies all external services except for Squid, forcing all World Wide Web connections to be established by the proxy.

If it is a firewall configuration including a DMZ, set the proxy there. In this case, it is important that all computers in the DMZ send their log files to hosts inside the secured network.

One way to implement this feature is with the aid of a “transparent” proxy. It will be covered in Section B on page 178.

Multiple Caches

“Multiple caches” means configuring different caches so that objects can be exchanged between them, reducing the total system load as well as increasing the chances of finding an object already in the local network. It enables the configuration of cache hierarchies so a cache is able to forward object requests to sibling caches or to a parent cache. It can get objects from another cache in the local network or directly from the source.

Choosing the appropriate topology for the cache hierarchy is very important, because it should not increase the overall traffic on the network. For example, in a very large network, it is possible to configure a proxy server for every subnetwork and connect it to a parent proxy, connected in its turn to the proxy cache from the ISP.

All this communication is handled by ICP (Internet Cache Protocol) running on top of the UDP protocol. Data transfers between caches are handled using HTTP (Hyper Text Transmission Protocol) based on TCP, but for these kinds of connections, it is preferable to use faster and simpler protocols capable of reacting to incoming requests within a maximum of one or two seconds.

To find the most appropriate server from which to get the objects, one cache sends an ICP request to all sibling proxies. These will answer the requests via ICP responses with a HIT code if the object was detected or a MISS if it was not. If multiple HIT responses were found, the proxy server will decide which server to download depending on factors such as which cache sent the fastest answer or which one is closer. If no satisfactory responses have been sent, the request will be sent to the parent cache.

Tip

To avoid duplication of objects in different caches in our network, other ICP protocols are used such as CARP (Cache Array Routing Protocol) or HTCP (HyperText Cache Protocol). The more objects maintained in the network, the greater the chances of finding the one we want.

Tip

Caching Internet Objects

Not all objects available in the network are static. There are a lot of dynamically generated CGI pages, visitor counters, or encrypted SSL content documents. This is the reason not to cache any object like this: every time you access one of this objects, it will already have changed again.

The question remains as to how long all the other objects stored in the cache should stay there. To determine this, all objects in the cache are assigned three different states:

1. **FRESH:** When this object is requested, it is sent without comparing it to the the original object on the web to see if it has changed.
2. **NORMAL:** The original server is queried to see if the object has changed. If it changed, the cache copy is updated.
3. **STALE:** The object is no longer considered valid and will be downloaded again from the server.

Web and proxy servers find out the status of an object by adding headers to these objects such as “Last modified” or “Expires” and the corresponding date. Other headers specifying that objects must not be cached are used as well.

Objects in the cache are normally replaced, due to a lack of free hard disk space, using algorithms such as LRU (Last Recently Used), which serve to replace cache objects. It consists of first replacing the less requested objects.

System Requirements

The most important thing is to determine the maximum load the system will have to bear. It is, therefore, important to pay more attention to the load picks, because these might be more than four times the day’s average. When in doubt, it would be better to overestimate the system’s requirements, because having Squid working close to the limit of its capabilities could lead to a severe loss in the quality of the service.

Hard Disk

Speed: choosing fast hard disks

Speed plays an important role in the caching process, so should be of utmost concern. In hard disks, this parameter is described as “random-seek time”, measured in milliseconds. As a rule of thumb, the lower this value, the better.

According to the Squid User’s Guide (<http://www.squid-cache.org>), for a system using only one disk, the formula for calculating the number of requests per second from the seek time of the disks is quite easy:

$$\text{requests per second} = 1000 / \text{seek time}$$

Squid enables more disks to be used simultaneously, increasing the number of requests per second. For instance, if you have three disks with the same seek time of 12 milliseconds, using the following formula will result in:

$$\begin{aligned} \text{requests per second} &= 1000 / (\text{seek time} / \text{number of disks}) \\ &= 1000 / (12/3) \\ &= 250 \text{ requests per second} \end{aligned}$$

In comparison to using IDE or SCSI disks, SCSI is preferable. Newer IDE disks, however, have similar seek times as SCSI and, together with DMA-compatible IDE controllers, increase the speed of data transfer without considerably increasing the system load.

Size of the Disk Cache

It depends on a few factors. In a small cache, the probability of a HIT (finding the requested object already located there) will be small, because the cache is easily filled up so the less requested objects will be replaced by newer ones. On the other hand, if 1 GB is available for the cache and the users only surf 10 MB a day, it will take more than 100 days to fill the cache.

Probably the easiest way to determine the needed cache size is to consider the maximum transfer rate of our connection. With a 1 MB/s connection, the maximum transfer rate will be 125 KB/s. If all this traffic ends up in the cache, in one hour it will add up to 450 MB and, assuming that all this traffic is generated in only 8 working hours, it will reach 3.6 GB in one day. Because the connection was not used up to its maximum capacity (otherwise we would have procured a faster one), we could assume that the total amount of data going through the cache is about 2 GB. In the example, to keep all the browsed data of *one* day in the cache, we will require 2 GB of disk space for Squid.

Summing up, Squid tends to read and write smaller blocks from or to the disk, making it more important how fast it detects these objects on the disk than having a fast disk.

RAM

The amount of memory required by Squid directly correlates to the amount of objects allocated in the cache. Squid also stores cache object references and frequently requested objects in memory to speed up the retrieval of this data. The memory is one million times faster than a hard disk. (Compare the seek

time of a hard disk, about 10 milliseconds, with the 10 nanoseconds access time of the newer RAM memories)

Every object in RAM memory has a size of 72 bytes (for “small” pointer architectures like Intel, Sparc, or MIPS. For Alpha, it is 104 bytes). If the average size of an object on the Internet is about 8 KB and we have 1 GB disk for the cache, we will be storing about 130,000 objects, resulting in close to 10 MB RAM only for meta data.

Squid also holds data in memory for lots of other stuff, such as a table with all used IP addresses, a fully qualified domain name cache, hot objects (the most requested), buffers, and access control lists.

It is very important to have more than enough memory for the Squid process, because, if it has to be swapped, the system performance will be dramatically reduced.

CPU

Squid is not a program that requires intensive CPU usage. The load of the processor is only increased while the contents of the cache are being loaded or checked. Using a multiprocessor machine does not increase the performance of the system. To increase efficiency, it is better to buy faster disks or add more memory.

Some examples of configured systems running Squid are available at <http://wwwcache.ja.net/servers/squids.html>.

Starting Squid

Squid is already preconfigured in SuSE Linux, so you can start it easily right after installation. A prerequisite for a smooth start is an already configured network, at least one name server and, of course, Internet access. Problems can arise if a dial-up connection is used with dynamic DNS configuration. In cases such as this, at least the name server should be entered clearly, because Squid will not start if it does not detect a DNS in the `/etc/resolv.conf`.

To start Squid, enter `rcsquid start` at the command line at root. For the initial start-up, the directory structure will first have to be defined in `/var/squid/cache`. This is done by the start script `/etc/init.d/squid` automatically and can take a few seconds or even minutes. If done appears to the right in green, Squid has been successfully loaded. Test Squid’s functionality on the local system by entering `localhost` and `Port 3128` as

proxy in the browser. To allow all users to access Squid and thus the Internet, change the entry in the configuration file `/etc/squid.conf` from `http_access deny all` to `http_access allow all`. However, in doing so, consider that Squid is made completely accessible to anyone by this action. Therefore, you should, in any case, define ACLs to control access to the proxy. More on this is available in Section B on page 176.

If you have made changes in the configuration file `/etc/squid.conf`, instruct Squid to load the changed file. Do this by entering `rcsquid reload` or restart Squid with `rcsquid restart`. Also, the command `rcsquid status` is important. With it, determine whether the proxy is running. With `rcsquid stop`, halt Squid. The latter can take a while, since Squid waits up to half a minute (`shutdown_lifetime`) before dropping the connections to the clients then will write its data to the disk. If Squid is halted with `kill` or `killall`, this can lead to the destruction of the cache, which will then have to be fully removed to restart Squid.

If Squid dies after a short period of time, although it has seemingly been started successfully, it can be the result of a faulty name server entry or a missing `/etc/resolv.conf` file. The cause of the start failure would then be logged by Squid in the `/var/squid/logs/cache.log` file.

If Squid should be loaded automatically when the system boots, reset the entry `START_SQUID=no` to `START_SQUID=yes` in the `/etc/rc.config` file.

An uninstall of Squid will neither remove the cache or the log files. Manually delete the `/var/squid` directory.

The Configuration File `/etc/squid.conf`

All Squid proxy server settings are made in the `/etc/squid.conf` file. To start Squid for the first time, no changes will be necessary in this file, but external clients will initially be denied access. The proxy needs to be made available for the localhost, usually with 3128 as port. The options are extensive and therefore provided with ample documentation and examples in the preinstalled `/etc/squid.conf` file. Nearly all entries begin with a ``#`` sign (the lines are commented out) and the relevant specifications can be found at the end of the file. The given values almost always correlate with the default values, so removing the comment signs without changing any of the parameters actually has little effect in most cases. It is better to leave the sample as it is and reinsert the options along with the modified parameters in the line below. In this way, easily interpret the default values and the changes.

If you have updated an earlier Squid version, it is recommended to edit the new `/etc/squid.conf` and only apply the changes made in the previous file. If you try to implement the old `squid.conf` again, you are running a risk that the configuration will no longer function, because options are always being modified and new changes added.

General Configuration Options

http_port 3128 This is the port where Squid listens for client requests. The default port is 3128, but 8080 is also common. You have the option here of specifying several port numbers separated by blank spaces.

cache_peer <hostname> <type> <proxy-port> <icp-port> Here, you can enter a parent proxy as “parent”, for example, or use that of the provider. As <hostname>, the name and IP address of the proxy to use are entered and, as <type>, parent. For <proxy-port>, the port number is to be entered which is also specified by the operator of the parent for use in the browser, usually 8080. Set the <icp-port> to 7 or 0 if the ICP port of the parent is not known and its use is irrelevant to the provider. In addition, default and no-query should be specified after the port numbers to strictly prohibit the use of the ICP protocol. Squid will then behave like a normal browser as far as the provider’s proxy is concerned.

cache_mem 8 MB This entry defines the maximum amount of disk space Squid can use for the caches. The default is 8 MB.

cache_dir ufs /var/squid/cache 100 16 256 The entry `cache_dir` defines the directory where all the objects are stored on disk. The numbers at the end indicate the maximum disk space in MB to use, as well as the number of directories in the first and second level. The `ufs` parameter should be left alone. The default is 100 MB occupied disk space in the `/var/squid/cache` directory and to create 16 subdirectories inside it which each contain 256 more subdirectories. When specifying the disk space to use, always leave sufficient reserve disk space. Values from a minimum of fifty to a maximum of eighty percent of the available disk space make the most sense here. The last two numbers for the directories should only be increased with caution, because too many directories can also lead to performance problems. If you have several disks that share the cache, enter several `cache_dir` lines.

cache_access_log /var/squid/logs/access.log path for log message

cache_log /var/squid/logs/cache.log path for log message

cache_store_log `/var/squid/logs/store.log` path for log message

These three entries specify the path where Squid will log all of its actions. Normally, nothing is changed here. If Squid is experiencing a heavy usage burden, it might make sense to distribute the cache and log files over several disks.

emulate_httptd_log off If the entry is set to `on`, obtain readable log files. Some evaluation programs cannot interpret this, however.

client_netmask `255.255.255.255` With this entry, mask the logged IP addresses in the log files to hide the clients' identity. The last digit of the IP address will be set to zero if you enter `255.255.255.0` here.

ftp_user `Squid@` With this, set the password which Squid should use for the anonymous FTP login. The login `anonymous` and your e-mail address as password are generally used to access public FTP servers, which saves the trouble of entering your user name and password each time you download FTP. `Squid@` without the domain is the default, because the clients can originate from any domain. It can still make sense, however, to specify a valid e-mail address here, because some FTP servers can check these for validity.

cache_mgr `webmaster` An e-mail address to which Squid sends a message if it unexpectedly crashes. The default is `webmaster`.

logfile_rotate `0` If you call up `squid -k rotate`, Squid can rotate secured log files. The files will be enumerated in this process and after reaching the specified value, the oldest file at that point will be overwritten. This value here normally is `0`, because archiving and deleting log files in SuSE Linux is carried out by a cronjob in the configuration file `/etc/logfiles`. The period of time after which the files are deleted is defined in the `/etc/rc.config` file via the `MAX_DAYS_FOR_LOG_FILES` entry.

append_domain `<domain>` With `append_domain`, specify which domain will automatically be appended when none is given. Usually, your own domain is entered here, so entering `www` in the browser suffices to guarantee access to your own web server.

forwarded_for `on` If you set the entry to `off`, Squid will remove the IP address and the system name of the client from the HTTP requests.

negative_ttl `5 minutes`; **negative_dns_ttl** `5 minutes`

Normally, you do not need to change these values. If you have a dial-up connection, however, the Internet may, at times, not be accessible.

Squid will make a note of the failed requests then refuse to issue new ones, although the Internet connection has been reestablished. In a case such as this, change the `minutes` to `seconds` then, after clicking on `Reload` in the browser, the dial-up process should be reengaged after a few seconds.

never_direct allow <acl_name> To prevent Squid from taking requests directly from the Internet, use this command to force connection to another proxy. You need to have previously entered this in `cache_peer`. If `all` is specified as the `<acl_name>`, force all requests to be forwarded directly to the parent. This might be necessary, for example, if you are using a provider which strictly stipulates the use of its proxies or denies its firewall direct Internet access.

Options for Access Controls

Squid provides an intelligent system that controls access to the proxy. By implementing ACLs, it can be configured easily and comprehensively. This involves lists with rules processed sequentially. ACLs must be defined before they can be used. Some default ACLs, such as `all` and `localhost`, already exist. After defining an ACL, implement it, for example, in conjunction with `http_access`.

acl <acl_name> <type> <data> An ACL requires at least three specifications to define it. The name `<acl_name>` can be arbitrarily chosen. For `<type>`, select from a variety of different options in the `ACCESS CONTROLS` section in the `/etc/squid.conf` file. The specification for `<data>` depends on the individual ACL type and can also be read from a file, for example, via host names, IP addresses, or URLs. In the following are some simple examples:

```
acl mysurfers srcdomain .my-domain.com acl teachers
src 192.168.1.0/255.255.255.0 acl students src
192.168.7.0-192.168.9.0/255.255.255.0 acl lunch time MTWHF
12:00-15:00
```

http_access allow <acl_name> `http_access` defines who is allowed to use the proxy and also who can access what on the Internet. For this, ACLs must be given. `localhost` and `all` have already been defined above, which can deny or allow access via `deny` or `allow`. A list containing any number of `http_access` entries can be created. They will be processed from top to bottom and, depending on which occurs first, access will be allowed or denied to the respective URL. The last entry

should always be `http_access deny all`. In the following example, the `localhost` has free access to everything while all other hosts are denied access completely.

```
http_access allow localhost http_access deny all
```

Another example, where the previously defined ACLs are used: The group `teachers` always has access to the Internet, while the group `students` only gets access Monday to Friday during lunch time.

```
http_access deny localhost http_access allow teachers
http_access allow students lunch time http_access deny all
```

The list with the `http_access` entries should only be entered, for the sake of readability, at the designated position in the `/etc/squid.conf` file — between the text

```
# INSERT YOUR OWN RULE(S) HERE TO ALLOW ACCESS FROM YOUR
CLIENTS
```

and the last

```
http_access deny all
```

redirect_program /usr/bin/squidGuard With this option, a redirector, such as `SquidGuard`, which is able to block unwanted URLs, can be specified. Internet access can be individually controlled for various user groups with the help of proxy authentication and the appropriate ACLs. `SquidGuard` is a package in and of itself which can be separately installed and configured.

authenticate_program /usr/sbin/pam_auth If users must be authenticated on the proxy, a corresponding program, such as `pam_auth`, can be specified here. When accessing `pam_auth` for the first time, the user will see a login window where the user name and password must be entered. In addition, an ACL is still required so only clients with a valid login can surf the Internet:

```
acl password proxy_auth REQUIRED
http_access allow password http_access deny all
```

The `REQUIRED` after `proxy_auth` can be replaced with a list of permitted user names or with the path to such a list.

ident_lookup_access allow *<acl_name>* With this, you will manage to have an ident request run through for all ACL-defined clients to find out each user's identity. If you apply `all` to the *<acl_name>*, this will be valid for all clients. Also, an ident daemon must be running on all clients. For Linux, install the `pidentd` package for this purpose. For Windows, there is free software available to download from the Internet. To ensure that only clients with a successful ident lookup are permitted, a corresponding ACL will also have to be defined here:

```
acl identhosts ident REQUIRED
http_access allow identhosts http_access deny all
```

Here, too, replace the `REQUIRED` with a list of permitted user names. Using `ident` can slow down the access time quite a bit, because ident lookups will definitely be repeated for each request.

Transparent Proxy Configuration

The usual way of working with proxy servers is the following: the web browser sends requests to a certain port in the proxy server and the proxy provides these required objects, whether they are in its cache or not. When working in a real network, several situations may arise:

- For security reasons, it is recommended that all clients use a proxy to surf the Internet.
- All clients must use a proxy whether they are aware of it or not.
- In larger networks already using a proxy, it is possible to spare yourself the trouble of reconfiguring each machine whenever changes are made in the system.

In all these cases, a transparent proxy may be used. The principle is very easy: the proxy intercepts and answers the requests of the web browser, so that the web browser receives the requested pages without knowing where they are coming from. This entire process is done transparently, hence the name.

Kernel Configuration

First, make sure that the proxy server's kernel has support for transparent proxies. Otherwise, add this option to the kernel and compile it again.

In the entry corresponding to Networking Options, select 'Network Firewalls' then the options 'IP: firewalling' and 'IP: Transparent proxying'. Now, save the new configuration, compile the new kernel, install it, reconfigure LILO if necessary, and restart the system.

Configuration Options in `/etc/squid.conf`

The options that need to be activated in the `/etc/squid.conf` file to get the transparent proxy up and running are:

- `httpd_accel_host` virtual
- `httpd_accel_port` port 80, where the actual HTTP server listens
- `httpd_accel_with_proxy` on
- `httpd_accel_uses_host_header` on

The default configuration as defined by the file `/etc/squid.conf` gives access to the proxy only from localhost. Therefore, it may be necessary to define additional access rules. See Section B on page 176.

Squid and Other Programs

In the following section, we will see how other applications interact with Squid. `cachemgr.cgi` enables the system administrator to check the amount of memory needed for caching objects, `squidgrd` filters web pages, and Calamaris is a report generator for Squid.

SquidGuard

This section is not intended to go through an extensive configuration of SquidGuard, only to introduce it and give some advice on using it. For more in-depth configuration issues, refer to the SquidGuard web site at <http://www.squidguard.org>

SquidGuard is a free (GPL), flexible, and ultra fast filter, redirector, and “access controller plugin” for Squid. It lets you define multiple access rules with different restrictions for different user groups on a Squid cache. SquidGuard uses Squid’s standard redirector interface.

SquidGuard can be used for the following:

- limit the web access for some users to a list of accepted or well known web servers or URLs.
- block access to some listed or blacklisted web servers or URLs for some users.
- block access to URLs matching a list of regular expressions or words for some users.
- redirect blocked URLs to an “intelligent” CGI-based info page.
- redirect unregistered users to a registration form.
- redirect banners to an empty GIF.
- have different access rules based on time of day, day of the week, date, etc.
- have different rules for different user groups.
- and much more

Neither SquidGuard or Squid can be used to:

- Edit, filter, or censor text inside documents
- Edit, filter, or censor HTML-embedded script languages such as JavaScript or VBScript

Using SquidGuard

Install the package `squidgrd` from the series `n`. Edit a minimal configuration file `/etc/squidguard.conf`. There are plenty of configuration examples in <http://www.squidguard.org/config/>. Experiment later with more complicated configuration settings.

The following step is to create a dummy “access denied” page or a more or less intelligent CGI page to redirect Squid in case the client requests a blacklisted web site. Again, using Apache is strongly recommended.

Now, tell Squid to use SquidGuard. Use the following entries in the `/etc/squid.conf` file:

```
redirect_program /usr/bin/squidGuard
```

There is another option called `redirect_children` configuring how many different “redirect” (in this case SquidGuard) processes are running on the machine. SquidGuard is fast enough to cope with lots of requests (SquidGuard is quite fast: 100,000 requests within 10 seconds on a 500MHz Pentium with 5900 domains, 7880 URLs, 13780 in sum). Therefore, it is not recommended to set more than 5 processes, because this may lead to an unnecessary increase of memory for the allocation of these processes.

```
redirect_children 5
```

Last of all, send a HUP signal to Squid to have it read the new configuration:

```
rcsquid reload
```

Test your settings with a browser.

Cache Report Generation with Calamaris

Calamaris is a Perl script used to generate reports of cache activity in ASCII or HTML format. It works with native Squid access log files. The Calamaris Home Page is located at <http://Calamaris.Cord.de/>

The use of the program is quite easy. Log in as `root`, then:

```
cat access.log.files | calamaris [options] > reportfile
```

It is important when piping more than one log file that the log files are chronologically ordered, with older files first.

The various options:

- a** normally used for the output of available reports
- w** an HTML report
- l** a message or logo in the header of the report

Further information on the various options can be found in the manual page `man calamaris`.

A typical example:

```
cat access.log.2 access.log.1 access.log | calamaris -a -w > /usr/local/httpd/htdocs/Squid/squidreport.html
```

We place the report in the directory of the web server. Again, Apache is required to view the reports.

Another powerful cache report generator tool is SARG (Squid Analysis Report Generator), included in series n. Further information on this can be found in the relevant Internet pages at <http://web.onda.com.br/orso/>

More Information on Squid

Visit the home page of Squid: <http://www.squid-cache.org/>. Here, find the Squid User Guide and a very extensive collection of FAQs on Squid.

The Mini-Howto regarding transparent proxies in the package howtoen, under `/usr/share/doc/howto/en/mini/TransparentProxy.gz`

In addition, mailing lists are available for Squid at:
squid-users@squid-cache.org.

The archive for this is located at:
<http://www.squid-cache.org/mail-archive/squid-users/>

Network Security

This chapter provides detailed information about several aspects of network security. It begins with information about masquerading. SSH is a protocol for remote logins over an encrypted connection. The last section provides more detailed information about general security issues.

Masquerading and Firewalls	184
SSH — Secure Shell, the Safe Alternative	190
Security and Confidentiality	195

Masquerading and Firewalls

Owing to its outstanding network capabilities, Linux is becoming more widespread as a router operating system for dial-up or dedicated lines. “Router,” in this case, refers to a host which has more than one network interface and transmits any packets not destined for one of its own network interfaces to another host communicating with it. This router is often called a gateway. The packet filtering mechanism provided by the Linux kernel allows precise control over which packets of the overall traffic are allowed through.

In general, defining the exact rules for a packet filter requires at least some experience on the part of the administrator. SuSEfirewall is highly configurable, making it a good choice for a more complex packet filtering setup. A Linux machine can be used as a router with masquerading to link a local network through a dial-up or dedicated connection where only one IP address is visible to the outside world. Masquerading is accomplished by implementing rules for packet filtering.

Caution

This chapter only describes standard procedures which should work well in most situations. However, there is no guarantee that this book or other materials provided by us are free from errors which might have escaped our attention.

Caution

Masquerading Basics

Masquerading is the Linux specific form of NAT (Network Address Translation). The basic principle is not very complicated: Your router has more than one network interface, typically a network card and a modem (or an ISDN interface). While one of these interfaces will link you with the outside world, the remaining ones are used to connect this router with the other hosts in your network. For example, the dial-up is conducted via ISDN and the network interface is `ipp0`. Several hosts in your local network are connected to the network card of your Linux router, in this example, `eth0`. The network address of the internal network is `192.168.0.0`, the router’s address is `192.168.0.1`, and the hosts connected to it have addresses like `192.168.0.2` and `192.168.0.3`. These hosts will send any packets not destined for the local network to the address `192.168.0.1`, the network interface of your default router or gateway.

Note

Make sure that both the broadcast addresses and the network masks are the same for all the hosts when configuring your network.

Note

As soon as one of the hosts sends a packet destined for an Internet address, this packet is sent to the network's default router. The router needs to be configured to actually forward such packets. SuSE Linux does not enable this with a default installation for security reasons. Set the variable `IP_FORWARD`, defined in the file `/etc/rc.config`, to `IP_FORWARD=yes`. The forwarding mechanism is enabled after rebooting or issuing this command:

```
earth:~ # echo 1 > /proc/sys/net/ipv4/ip_forward
```

This is where masquerading begins. The router has only one IP address visible from the outside (in our example the address of the ISDN interface after dial-up). The source address of an outgoing packet must be replaced with the router's own address before sending it out over the external network interface. If the router did not replace the source address, the receiving end would have no means to reply. This is especially the case if you are using the `192.168.x.x` address range. Although it represents a valid set of IP addresses, they are not forwarded at all by any of the Internet's routers.

The target host at the other end of the link only knows your router, but not the host in your internal network that sent the packet. Your internal host disguises itself behind the router, which is why the technique is called "masquerading". The router will consequently be the destination of any reply packets. Therefore, it has to identify the incoming packets, change the target address to the intended recipient, and forward it to that host in the local network.

The identification of packets belonging to a connection handled by a masquerading router is done with the help of a table that is kept in the kernel of your router as long as the connection is active. By using the `ipchains` and the `iptables` commands, the superuser (`root`) can view these tables. Read the man pages for these commands for detailed instructions. For the identification of single masqueraded connections not only the source and target addresses are relevant, but also the port numbers and the protocols involved. With this method, your router is capable of hiding many thousand connections per internal host simultaneously.

With the routing of inbound traffic depending on the masquerading table, there is no way to open a connection to some internal host from the outside.

For such a connection, there would be no entry in the table because, the entry itself is only created if an internal host opens a connection with the outside. In addition, any established connection is assigned a status entry in the table and this entry cannot be used by another connection. A second connection would require another status record.

As a consequence of all this, you might experience some problems with a number of applications: programs use protocols to talk to each other and some of these will try to open additional connections or send packets from the server to your client which cannot be recognized by a simple packet filter as being valid. Examples of such protocols are ICQ, cucme, IRC (DCC, CTCP), Quake, and FTP (in PORT mode). Netscape, as well as the standard ftp program and many others, uses the PASV mode. This passive mode is much less problematic as far as packet filtering and masquerading is concerned. The FTP protocol opens a controlling connection in addition to the data connection for the file transfer. In PORT mode, the server opens a connection to the client. In PASV (passive) mode, the client establishes a connection. As stated previously, our setup allows for connections to be opened exclusively from the internal side, which explains the trouble that FTP will cause if used in PORT mode.

Firewalling Basics

“Firewall” is probably the most widely used term to describe a mechanism to link two networks and control the data traffic between them. There are various types of firewalls which mostly differ in regard to the abstract level on which traffic is analyzed and controlled. Strictly speaking, the mechanism described in this section is called a “packet filter.” Like any other type of firewall, a packet filter alone does not guarantee full protection from all security risks. What a packet filter does is implement a set of rules related to protocols, ports, and IP addresses to decide whether data may pass through. This blocks any packets that, according to their addresses, are not supposed to reach your network. Packets sent to the telnet service of your hosts on port 23, for example, should be blocked, while you might want people to have access to your web server and therefore enable the corresponding port. Note that a packet filter will not scan the contents of any packets as long as they have legitimate addresses (e.g., directed to your web server). Thus, packets could attack your CGI server, but the packet filter would let them through.

A more effective, but also more complex, mechanism is the combination of several types of systems, such as a packet filter interacting with an application gateway or proxy. In this case, the packet filter rejects any packets destined to disabled ports. Only packets directed to the application gateway are

allowed through. This gateway or proxy pretends to be the actual client of the server. In a sense, such a proxy could be considered a masquerading host on the protocol level used by the application. One example for such a proxy is Squid, an HTTP proxy server. To use Squid, the browser needs to be configured to communicate via the proxy, so that any HTTP pages requested would be served from the proxy cache rather than directly from the Internet. As another example, the SuSE proxy suite (the package proxy-suite in series sec) includes a proxy for the FTP protocol.

SuSEfirewall

The SuSEfirewall is a script used for protecting the adminhost. This section describes the configuration of SuSEfirewall, a rather more challenging task. It requires a certain degree of experience and understanding. Find documentation about SuSEfirewall in `/usr/share/doc/packages/SuSEfirewall`. The theoretical background is also covered in this manual, see Chapter C on page 195.

The configuration of SuSEfirewall is stored in the file `/etc/rc.config.d/firewall.rc.config` and is commented in English. In the following we demonstrate a successful configuration step by step. For each configuration item, find a note as to whether it is relevant for firewalling or masquerading. If you stumble across any comments in the configuration file that are related to what is called DMZ (or “demilitarised zone”), this is not covered here.

If your requirements are strictly limited to masquerading, fill out the items marked with *masquerading* only.

- **START_FW** (firewall, masquerading): Set this variable to `yes` in `/etc/rc.config`, to ensure that the script is started. This enables the firewall or masquerading.
- **FW_DEV_WORLD** (firewall, masquerading): For example, `eth0` as the device linked to the Internet. In the case of ISDN, choose `ipp0` here.
- **FW_DEV_INT** (firewall, masquerading): The device linking you with the internal, “private” network. Leave this blank if there is no internal network and the firewall is supposed to protect only the one host.
- **FW_ROUTE** (firewall, masquerading): If you need the masquerading function, enter `yes` here. Your internal hosts will not be invisible to the outside, because their private network addresses (e.g. `192.168.x.x`) are ignored by Internet routers.

For a firewall without masquerading, only set this to *yes* if you want to allow access to the internal network. Your internal hosts need to use officially registered IPs in this case. Normally, however, you should *not* allow access to your internal network from the outside.

- **FW_MASQUERADE** (masquerading): Set this to *yes* if you need the masquerading function. Note that it is more secure to have a proxy server between the hosts of the internal network and the Internet.
- **FW_MASQ_NETS** (masquerading): Specify the hosts or networks to be masqueraded, leaving a space between the individual entries. For example:

```
FW_MASQ_NETS="192.168.0.0/24 192.168.10.1"
```

- **FW_PROTECT_FROM_INTERNAL** (firewall): Set this to *yes* to protect your firewall host from attacks originating in your internal network. Services will only be available to the internal network if explicitly enabled. See also **FW_SERVICES_INTERNAL_TCP** and **FW_SERVICES_INTERNAL_UDP**.
- **FW_AUTOPROTECT_GLOBAL_SERVICES** (firewall): This should normally be *yes*.
- **FW_SERVICES_EXTERNAL_TCP** (firewall): Enter the services that should be available, e.g., "www smtp ftp domain 443" . Normally leave this blank for a workstation at home that is not intended to offer any services.
- **FW_SERVICES_EXTERNAL_UDP** (firewall): Leave this blank if you do not run a name service that you want to make available to the outside. Otherwise, enter the ports to be used.
- **FW_SERVICES_INTERNAL_TCP** (firewall): This defines the services available to the internal network. The notation is the same as for external TCP services, but, in this case, refers to the *internal* network.
- **FW_SERVICES_INTERNAL_UDP** (firewall): See above.
- **FW_TRUSTED_NETS** (firewall): Specify the hosts you *really* trust ("trusted hosts"). Note, however, that these need to be protected from attacks, too.

For example: "172.20.0.0/16 172.30.4.2" means that all hosts which have an IP address beginning with 172.20.x.x, along with the host with the IP address 172.30.4.2, are allowed to pass through the firewall.

- **FW_SERVICES_TRUSTED_TCP** (firewall): Here, specify the port addresses which may be used by the "trusted hosts". For example, to grant them access to all services, enter 1:65535. Usually, it is sufficient to enter `ssh` as the only service.
- **FW_SERVICES_TRUSTED_UDP** (firewall): Just like above, but for UDP ports.
- **FW_ALLOW_INCOMING_HIGHPORTS_TCP** (firewall): Set this to `ftp-data` if you intend to use normal (active) FTP services.
- **FW_ALLOW_INCOMING_HIGHPORTS_UDP** (firewall): Set this to `dns` to use the name servers registered in `/etc/resolv.conf`. If you enter `yes` here, all high ports will be enabled.
- **FW_SERVICE_DNS** (firewall): Enter `yes` if you run a name server that is supposed to be available to external hosts. At the same time, enable port 53 under `FW_TCP_SERVICES_*`.
- **FW_SERVICE_DHCLIENT** (firewall): Enter `yes` here if you use `dhclient` to get your IP address assigned.
- **FW_LOG_*** (firewall): Specify the firewall's logging activity. For normal operation, it is sufficient to set `FW_LOG_DENY_CRIT` to `yes`.
- **FW_STOP_KEEP_ROUTING_STATE** (firewall): Insert `yes` if you have configured your dial-up procedure to work automatically via `ciold` or ISDN (dial on demand).

Now that you have configured SuSEfirewall, do not forget to test your setup (for example, with `telnet` from an external host). Have a look at `/var/log/messages`, where you should see something like:

```
Feb  7 01:54:14 www kernel: Packet log: input DENY eth0
PROTO=6 129.27.43.9:1427 195.58.178.210:23 L=60 S=0x00
I=36981 F=0x4000 T=59 SYN (#119)
```

SSH — Secure Shell, the Safe Alternative

In these times of increasing networks, accessing a remote system also becomes more common. Regardless of the activity, the person accessing the system must be authenticated.

Most users should know by now that the user name and password are only intended for individual use. Strict confidence pertaining to personal data is usually guaranteed between the employer, computer center, or service provider. However, the ongoing practice of authenticating and transferring data in clear text form is a frightening phenomenon. Most directly affected are the commonly used services Post Office Protocol (POP) for retrieving mail and telnet for logging in on remote systems. Using these methods, user information and data considered sensitive, such as the contents of a letter or a chat via the talk command, travel openly and unsecured over the network. This encroaches on the user's privacy and leaves such access methods open to misuse. Usually, this misuse occurs by accessing one system to attack another or to obtain administrator or root permissions.

Any device involved in data transfer or operating on the local network, such as firewall, router, switch, mail servers, or workstations, can also access the data. There are laws prohibiting such behavior, but it is difficult to detect.

The SSH software provides the necessary protection. Complete authentication, usually user name and password, as well as the communication is encrypted. Even here, snatching the transferred data is possible, but the contents cannot be deciphered by intruders without the key. This enables secure communication via unsafe networks, such as the Internet. SuSE Linux, provides the package OpenSSH.

The OpenSSH Package

SuSE Linux installs the package OpenSSH by default. The programs `ssh`, `scp`, and `sftp` will thus be available as alternatives to `telnet`, `rlogin`, `rsh`, `rcp`, and `ftp`.

The ssh Program

Using the `ssh` program, it is possible to log in to remote systems and work interactively. It replaces both `telnet` and `rlogin`. The symbolic name `slogin` points to `ssh`. For example, it is possible to log in to the host `sun` with the command `ssh sun`. The host then prompts for the password on `sun`.

Following successful authentication, work from the command line there or use interactive applications. If the local user name is different from the remote user name, log in using a different login name with `ssh -l augustine sun` or `ssh augustine@sun`.

Furthermore, `ssh` offers the option of running commands on another system, as does `rsh`. In the following example, we will run the command `uptime` on the host `sun` and create a directory with the name `tmp`. The program output will be displayed on the local terminal of the host `earth`.

```
newbie@earth:~ > ssh sun"uptime; mkdir tmp"
newbie@sun's password:
1:21pm up 2:17, 9 users, load average: 0.15, 0.04, 0.02
```

Quotation marks are necessary here to send both instructions with one command. It is only by doing this that the second command is likewise executed on `sun`.

scp — Secure Copy

`scp` copies files to a remote machine. It is the secure and encoded substitute for `rcp`. For example, `scp MyLetter.tex sun:` copies the file `MyLetter.tex` from the machine `earth` to the machine `sun`. To give a different user name, use the `username@machine` format.

After the correct password is entered, `scp` starts the data transfer and shows a series of stars, gradually marking the progress from left to right. In addition, the estimated time of arrival will be shown in the right margin. All output can be suppressed by giving the option `-q`.

`scp` also provides a recursive copying feature for entire directories.

`scp -r src/ sun:backup/` copies the entire contents of the directory `src/` including all subdirectories to the machine `sun` in the subdirectory `backup/`. If this subdirectory does not exist yet, it will be created automatically.

Via the option `-p`, `scp` leaves the time stamp of the files unchanged. `-C` compresses the data transfer. This minimizes the data volume to be transferred, but creates heavier burden on the processor.

sftp — Secure File Transfer

Instead of `scp`, `sftp` can be used for secure file transfer. During the session, `sftp` provides many of the commands used by `ftp`. This may be an advantage over `scp`, especially when transferring data for which the file names are unknown.

The SSH Daemon (sshd) — Server-Side

To work with the SSH client programs `ssh` and `scp`, a server, the SSH daemon, has to be running in the background. This waits for its connections on TCP/IP port 22.

The daemon generates three key pairs when starting for the first time. The key pairs consist of a private and a public key. Therefore, this procedure is referred to as public key-based. To guarantee the security of the communication via SSH, only the system administrator can see the private key files. The file permissions are restrictively defined by the default setting. The private keys are only required locally by the SSH daemon and must not be given to anyone else. The public key components (recognizable by the name extension `.pub`) are sent to the communication partner and are readable for all users.

A connection is initiated by the SSH client. The waiting SSH daemon and the requesting SSH client exchange identification data comparing the protocol and software versions and preventing connection to the wrong port. Since a child process of the original SSH daemon replies to the request, several SSH connections can be made simultaneously.

The SSH protocol is available in two versions, 1 and 2, for the communication between SSH server and SSH client.

When using SSH with version 1, the server will then send its public host key and a server key, regenerated by the SSH daemon every hour. Both allow the SSH client to encrypt a freely chosen session key then send it to the SSH server. The SSH client also tells the server which encryption method (cipher) to use.

SSH in version 2 does not require a server key. A Diffie-Helman algorithm is employed instead for exchanging the keys.

The private host and server keys absolutely necessary for decoding the session key cannot be derived from the public parts. Only the SSH daemon contacted can decipher the session key using its proprietary keys (see also `/usr/share/doc/packages/openssh/RFC.nroff`). This initial connection phase can be watched closely using the SSH client program's error search option `-v`. Version 2 of the SSH protocol is used by default, which however can be overridden to use version 1 of the protocol with the `-1` switch. By storing all public host keys after initial contact in `~/.ssh/known_hosts` on the client side, so-called "man-in-the-middle" access attempts can be prevented. SSH servers that try to fraudulently use names and IP addresses of others will be exposed by a clear indicator. They will either be noticed due to a wrong host key which differs from `~/.ssh/known_hosts` or they cannot decipher the session key in the absence of an appropriate private counterpart.

It is recommended to securely archive the private and public keys stored in `/etc/ssh/` externally. In this way, key modifications can be detected and the old ones can be used again after a new installation. This spares users the unsettling warning. If it is verified that, despite the warning, it is indeed the correct SSH server, the existing entry regarding this system will have to be removed from `~/.ssh/known_hosts`.

SSH Authentication Mechanisms

Now the actual authentication will take place, which, in its simplest form, consists of entering a password as mentioned above. The goal of SSH was to introduce a secure software that is also easy to use. As it is meant to replace `rsh` and `rlogin` programs, SSH must also be able to provide an authentication method good for daily use. SSH accomplishes this by way of another key pair generated by the user. The SSH package also provides a help program, `ssh-keygen`, for this. After entering `ssh-keygen -t rsa` or `ssh-keygen -t dsa`, the key pair will be generated and you will be prompted for the base file name in which to store the keys:

```
Enter file in which to save the key (/home/newbie/.ssh/id_rsa):
```

Confirm the default setting and answer the request for a passphrase. Even if the software suggests an empty passphrase, a text from ten to thirty characters is recommended for the procedure described here. Do not use short and simple words or phrases. Confirm by repeating the passphrase. Subsequently, you will see where the private and public keys are stored, in our example, the files `id_rsa` and `id_rsa.pub`.

```
Enter same passphrase again: Your identification has been
saved in /home/newbie/.ssh/id_rsa Your public key has been
saved in /home/newbie/.ssh/id_rsa.pub. The key fingerprint is:
79:c1:79:b2:e1:c8:20:c1:89:0f:99:94:a8:4e:da:e8 newbie@sun
```

Use `ssh-keygen -p -t rsa` or `ssh-keygen -p -t dsa` to change your old passphrase.

Copy the public key component (`id_rsa.pub` in our example) to the remote machine and save it there at the location `~/.ssh/authorized_keys2`. You will be asked to authenticate yourself with your passphrase the next time you establish a connection. If this does not occur, verify the location and contents of these files.

In the long run, this procedure is more troublesome than giving your password each time. Therefore, the SSH package provides another tool, the

`ssh-agent`, which retains the private keys for the duration of an X session. The entire X session will be started as a child process of `ssh-agents`. The easiest way to do this is to set the variable `usessh` at the beginning of the `.xsession` file to `yes` and log in via a display manager such as KDM or XDM. Alternatively, enter `ssh-agent startx`.

Now you can use `ssh` or `scp` as usual. If you have distributed your private key as described above, you are no longer prompted for your password. Take care of terminating your X session or locking it with a password-protection, for instance `xlock`.

All the relevant changes which resulted from the introduction of version 2 of the SSH protocol have also been documented in the file `/usr/share/doc/packages/openssh/README.SuSE`.

X, Authentication, and Other Forwarding Mechanisms

Beyond the previously described security-related improvements, `ssh` also simplifies the use of remote X applications. If you run `ssh` with the option `-X`, the `DISPLAY` variable will automatically be set on the remote machine and all X output will be exported to the remote machine over the existing `ssh` connection. At the same time, X applications started remotely and locally viewed with this method cannot be intercepted by unauthorized persons.

By adding the option `-A`, the `ssh-agent` authentication mechanism will be carried over to the next machine. This way, you can work from different machines without having to enter a password, but only if you have distributed your public key to the destination hosts and properly saved it there.

Both mechanisms are deactivated in the default settings, but can be permanently activated at any time in the system-wide configuration file `/etc/ssh/sshd_config` or the user's `~/.ssh/config`.

`ssh` can also be used to redirect TCP/IP connections. In the following example, the SMTP and POP3 port is redirected through `ssh`:
`ssh -L 25:sun:25 sun`. Here, each connection directed to "earth port 25", SMTP is redirected to the SMTP port on `sun` via an encrypted channel. This is especially useful for those using SMTP servers without SMTP-AUTH or POP-before-SMTP features. From any arbitrary location connected to a network, e-mail can be transferred to the "home" mail server for delivery. In a similar manner, the following command forwards all port 110 and POP3 requests on `earth` to the POP3 port of `sun`: `ssh -L 110:sun:110 sun`.

Both examples must be carried out by user `root`, because the connection is made to privileged local ports. E-mail is sent and retrieved by normal

users in an existing SSH connection. The SMTP and POP3 host must be set to `localhost` for this.

Additional information can be found in the manual pages for each of the programs described above and also in the files under `/usr/share/doc/packages/openssh`.

Security and Confidentiality

Basic Considerations

One of the main characteristics of a Linux or UNIX system is its ability to handle several users at the same time (multiuser) and to allow these users to perform several tasks (multitasking) on the same computer simultaneously. Moreover, the operating system is network transparent. The users often do not know whether the data or applications they are using is provided locally from their machine or made available over the network.

With the multiuser capability the respective data of different users must be stored separately. Security and privacy need to be guaranteed. “Data security” was already an important issue, even before computers could be linked through networks. Just like today, the most important concern was the ability to keep data available in spite of a lost or otherwise damaged data medium, a hard disk in most cases.

This chapter is primarily focused on confidentiality issues and on ways to protect the privacy of users, but it cannot be stressed enough that a comprehensive security concept should always include procedures to have a regularly updated, workable, and tested backup in place. Without this, you could have a very hard time getting your data back — not only in the case of some hardware defect, but also if the suspicion arises that someone has gained unauthorized access and tampered with files.

Local Security and Network Security

There are several ways of accessing data:

- personal communication with people who have the desired information or access to the data on a computer
- directly from the console of a computer (physical access)
- over a serial line

- using a network link

In all these cases, a user should be authenticated before accessing the resources or data in question. A web server might be less restrictive in this respect, but you still would not want it to disclose all your personal data to any surfer out there. On a SuSE system, a few tweaks are sufficient to make it boot right into your desktop without even asking for a password, but, in most cases, that would not be such a good idea, as anybody could change data or run programs.

In the list above, the first case is the one where the highest amount of human interaction is involved, such as when you are contacting a bank employee and are required to prove that you are the person owning that bank account. Then you will be asked to provide a signature, a PIN, or a password to prove that you are the person you claim to be. In some cases, it might be possible to elicit some intelligence from an informed person just by mentioning known bits and pieces here and there to win the confidence of that person by using clever rhetoric. The victim could be led to gradually reveal more information, maybe without even becoming aware of it.

Some people are rather unmindful of what they say or act unconsciously in the way they give answers, so that even a question which they believe was left unanswered might provide enough information to proceed with an even more precise question. Piece after piece gets added to the puzzle until the picture is nearly complete (“No, Mr. Smith is on vacation right now, it’s at least three weeks before he’ll be back in. He’s not my boss anyway, you know he’s up there in the fourth floor while I’m here in the third!”). Among hackers, this is called “social engineering”. You can only guard against this by educating people and by dealing with language and information in a conscious way. Before breaking into computer systems, attackers often try to target receptionists, service people working with the company, or even family members and, in many cases, such an attack based on social engineering will only be discovered at a much later time.

A person wanting to obtain unauthorized access to your data could also use the traditional way and try to get at your hardware directly. Therefore, the machine should be protected against any tampering so that no one can remove, replace, or cripple its components. This also applies to backups and even any network cable or the power cord. Likewise, it might be necessary to secure the boot procedure, as there are some well-known key combinations which invoke special reactions during booting. Protect yourself against this by setting passwords for the BIOS and the bootloader.

Serial terminals connected to serial ports are still used in many places, but are rarely installed with new systems anymore. With regard to data access,

serial terminals are a special case. Unlike network interfaces, they do not rely on a network protocol to communicate with the host. A simple cable, or maybe an infrared port, is used to send plain characters back and forth between the devices. The cable itself is the weakest point of such a system: with an older printer connected to it, it is really easy to record anything that runs over the wires. What can be achieved with a printer can also be accomplished in other ways, depending on the effort that goes into the attack.

Networks make it easier for us to access data remotely, but they do this with the help of network protocols which are often rather complex. This might seem paradoxical at first, but is really indispensable if you want to remotely control a computer or to retrieve data from it no matter where you are. It is necessary to have abstract, modular designs with layers that are more or less separate from each other. We rely on such modular designs in many daily computing situations. Modularity means that your text processor, for example, does not need to know about the kind of hard disk you use or your e-mail program should not be concerned with whether you have a modem or an ethernet card. Components of your operating system, Linux in this case, provide the necessary functions and make these available to the system through a predefined interface. With this modularity, a text processor or a mail user agent (MUA) can function on a variety of hardware platforms and you can run them from some place in the world with the necessary equipment.

Regarding the data, there is no difference between opening a file from a command line or looking at it with a web browser. The file could also be read via a network (using a telnet program or with a secure shell client — which is actually a much better option as ssh encrypts all network traffic). To do so, the host and the network need to be connected and the user needs to log in and authenticate. The possible actions are still restricted, however, by the file permissions.

Reading a file locally on a host requires other access rules than opening a network connection with a server on a different host. There is a distinction between local security and network security. The line is drawn where data has to be put into packets to be sent somewhere else.

Local Security

Local security starts with the physical environment in the location where the computer is running. Assume that your machine is set up in a place where security is in line with your expectations and needs.

The main goal of “local security” is to keep users separate from each other, so that no user can assume the permissions or the identity of another one.

This is a general rule to be observed, but it is especially true for the user `root` who holds the supreme power on the system. User `root` can take on the identity of any other local user without being prompted for the password and read any locally stored file.

For an attacker who has obtained access to local resources from the command line, there is certainly no shortage of things that could be done to compromise the system.

Passwords

On a Linux system, passwords are, of course, *not* stored as plain text and the text string entered is not simply matched with the saved pattern. If this were the case, all accounts on your system would be compromised as soon as someone got access to the corresponding file. Instead, the stored password is encrypted and, each time it is entered, is encrypted again and the two encrypted strings are compared. Naturally, this will only work if the encrypted password cannot be reverse-computed into the original text string. This is actually achieved by a special kind of algorithm, also called “trapdoor algorithm,” because it only works in one direction. An attacker who has obtained the encrypted string will not be able to get your password by simply applying the same algorithm again. Instead, it would be necessary to test all the possible character combinations until a combination is found which looks like your password when encrypted. As you can imagine, with passwords that are eight characters long, there are quite a number of possible combinations to calculate.

In the seventies, it was argued that this method would be more secure than others due to the relative slowness of the algorithm used, which took a few seconds to encrypt just one password. In the meantime, however, PCs have become powerful enough to do several hundred thousand or even millions of encryptions per second. Because of this, encrypted passwords should not be visible to regular users (`/etc/shadow` cannot be read by normal users). It is even more important that passwords are not easy to guess, in case the password file becomes visible due to some error. Consequently, it is not really useful to “translate” a password like “tantalise” into “t@nt@1ls3”.

Replacing some letters of a word with similar looking numbers is not safe enough. Password cracking programs which use dictionaries to guess words also play with substitutions like that. A better way is to make up a word with no common meaning, something which only makes sense to you personally, like the first letters of the words of a sentence or the title of a book, such as “The Name of the Rose” by Umberto Eco. This would give the following safe password: “TNotRbUE9”. By contrast, passwords like “beer-

buddy" or "jasmine76" are easily guessed even by someone who has only some casual knowledge about you.

The Boot Procedure

Configure your system so it cannot be booted from a floppy or from CD, either by removing the drives entirely or by setting a BIOS password and configuring the BIOS to allow booting from a hard disk only.

Normally, a Linux system will be started by a boot loader, allowing you to pass additional options to the booted kernel. This is crucial to your system's security. Not only does the kernel itself run with root permissions, but it is also the first authority to grant root permissions at system start-up. Prevent others from using such parameters during boot by using the options "restricted" and "password=your_own_password" in `/etc/lilo.conf`. Execute the command `lilo` after making any changes to `/etc/lilo.conf` and look for any unusual output the command might produce. If you forget this password, you will have to know the BIOS password and boot from CD to read the entry in `/etc/lilo.conf` from a rescue system.

File Permissions

As a general rule, always work with the most restrictive privileges possible for a given task. For example, it is definitely not necessary to be `root` to read or write e-mail. If the mail program you use has a bug, this bug could be exploited for an attack which will act with exactly the permissions of the program when it was started. By following the above rule, minimize the possible damage.

The permissions of the more than 200,000 files included in a SuSE distribution are carefully chosen. A system administrator who installs additional software or other files should take great care when doing so, especially when setting the permission bits. Experienced and security-conscious system administrators always use the `-l` option with the command `ls` to get an extensive file list, which allows them to detect any wrong file permissions immediately. An incorrect file attribute does not only mean that files could be changed or deleted. These modified files could be executed by `root` or, in the case of configuration files, that programs could use such files with the permissions of `root`. This significantly increases the possibilities of an attacker. Attacks like this are called cuckoo eggs, because the program (the egg) is executed (hatched) by a different user (bird), just like a cuckoo would trick other birds into hatching its eggs.

A SuSE Linux system includes the files `permissions`, `permissions.easy`, `permissions.secure`, and `permissions.paranoid`, all in the `/etc`

directory. The purpose of these files is to define special permissions, such as world-writable directories or, for files, the setuser ID bits, which means the corresponding program will not run with the permissions of the user that has launched it, but with the permissions of the file owner, root in most cases. An administrator may use the file `/etc/permissions.local` to add his own settings. The variable `PERMISSION_SECURITY`, set in `/etc/rc.config`, defines which of the above files is used by SuSE's configuration programs to set permissions accordingly. As a more convenient way to select the files, use the submenu 'Security' in YaST or YaST2. To learn more about the topic, read the comments in `/etc/permissions` or consult the manual page of `chmod` (`man chmod`).

File Race Conditions

Assume that a program wants to create a file in a directory which is world-writable (such as `/tmp`). First, the program checks whether the file already exists and, if that is not the case, creates it. However, between checking and file creation, there is a short moment which can be used by an attacker to create a symbolic link, a pointer to another file. The program may then be tricked into following the symbolic link, overwriting the target file with its own permissions. This is called a race because the interval during which the attacker can create a "symlink" is very short. The race is only possible if the checking and file creation procedure is not atomic (indivisible). If the race is allowed to take place at all, there is a chance that it may be won by the attacker. It is all a matter of probability.

Buffer Overflows and Format String Bugs

Special care must be taken whenever a program is supposed to process data which can or could be changed by a user, but this is more of an issue for the programmer of an application than for regular users. The programmer has to make sure that his application will interpret data in the correct way, without writing them into memory areas that are too small to hold them. Also, the program should hand over data in a consistent manner, using the interfaces defined for that purpose.

A "buffer overflow" can happen if the actual size of a memory buffer is not taken into account when writing to that buffer. There are cases where this data (as generated by the user) uses up some more space than what is available in the buffer. As a result, data is written beyond the end of that buffer area, which, under certain circumstances, makes it possible that a program will execute program sequences influenced by the user (and not by the programmer), rather than just processing user data. A bug of this kind may

have serious consequences, in particular if the program is being executed with special privileges (see Section C on page 199).

“Format string bugs” work in a slightly different way, but again it is the user input which could lead the program astray. In most cases, these programming errors are exploited with programs executed with special permissions — `setuid` and `setgid` programs — which also means that you can protect your data and your system from such bugs by removing the corresponding execution privileges from programs. Again, the best way is to apply a policy of using the lowest possible privileges (see Section C on page 199).

Given that buffer overflows and format string bugs are bugs related to the handling of user data, they are not only exploitable if access has been given to a local account. Many of the bugs that have been reported can also be exploited over a network link. Accordingly, buffer overflows and format string bugs should be classified as being relevant for both local and network security.

Viruses

Contrary to what some people will tell you, there *are* viruses that run on Linux. However, the viruses that are known were released by their authors as “proof of concept”, meaning that they were written to prove that the technique works as intended. On the other hand, none of these viruses have been spotted “in the wild” so far.

Viruses would not be able to survive and spread without a host on which they can live. In our case, the host would be a program or an important storage area of the system, such as the master boot record, which needs to be writable for the program code of the virus. Owing to its multiuser capability, Linux can restrict write access to certain files, which is the case especially with system files. Therefore, if you did your normal work with `root` permissions, you would increase the chance of the system being infected by a virus. By contrast, if you follow the principle of using the lowest possible privileges as mentioned above, chances of getting a virus are slim. Apart from that, you should never rush into executing a program from some Internet site that you do not really know. SuSE’s RPM packages carry a cryptographic signature as a digital label that the necessary care was taken to build them. Viruses are a typical sign that the administrator or the user lacks the required security awareness, putting at risk even a system that should be highly secure by its very design.

Viruses should not be confused with worms which belong to the world of networks entirely. Worms do not need a host to spread.

Network Security

Local security is concerned with keeping different users on one system apart from each other, especially from `root`. Network security, on the other hand, means that the system needs to be protected from an attack originating in the network.

The typical login procedure requiring a user name and a password for user authentication is a local security issue. However, in the particular case of logging in over a network, we need to differentiate between both security aspects. What happens until the actual authentication is network security and anything that happens afterwards is local security.

X Window System (X11 authentication)

As mentioned at the beginning, network transparency is one of the central characteristics of a UNIX system. X11, the windowing system of UNIX operating systems, can make use of this feature in an impressive way. With X11, it is basically no problem to log in at a remote host and start a graphical program that will then be sent over the network to be displayed on your computer. The protocol to communicate between the X application and the X server (which is the local process that draws the windows with the help of your video card) is relatively lightweight as far as bandwidth usage is concerned. This is because the protocol was designed in the eighties when network bandwidth was still a scarce resource.

Now if we want an X client to be displayed remotely using our X server, the latter is supposed to protect the resource managed by it (i. e. the display) from unauthorized access. In more concrete terms, certain permissions must be given to the client program. With the X Window System, there are two ways to do this, called host-based access control and cookie-based access control. The former relies on the IP address of the host where the client is supposed to run; the program to control this is `xhost`. What `xhost` does is to enter the IP address of a legitimate client into a tiny database belonging to the X server. Note, however, that relying on IP addresses for authentication is not very secure. For example, if there were a second user working on the host sending the client program, that user would have access to the X server as well — just like someone stealing the IP address. Because of these shortcomings, we will not describe this authentication method in more detail here, but you can learn about the way it functions if you read the man page of `xhost`, which includes a similar warning.

In the case of cookie-based access control, a character string is generated which is only known to the X server and to the legitimate user, just like an

ID card of some kind. This cookie (the word goes back not to ordinary cookies, but to Chinese fortune cookies which contain an epigram) is stored on login in the file `.xauthority` in the user's home directory and is available to any X Window client wanting to use the X server to display a window. The file `.xauthority` can be examined by the user with the tool `xauth`. If you were to rename `.xauthority` or if you deleted the file from your home directory by accident, you would not be able to open any new windows or X clients. Read more about X Window security mechanisms in the man page of `Xsecurity` (`man Xsecurity`).

Apart from that, `ssh` (secure shell) can be used to completely encrypt a network connection and forward it to an X server transparently without the encryption mechanism being perceived by the user. This is also called X forwarding. X forwarding is achieved by simulating an X server on the server side and setting a `DISPLAY` variable for the shell on the remote host. Before being displayed, the client opens a connection with `sshd` (secure shell daemon, the server side program), which then gets the connections through to the real X server. If your setup requires that X clients are displayed remotely, consider using `ssh`. The man page of `ssh` has more information about the functionality of this program. Further details about `ssh` can be found in Section C on page 190 of this book.

Caution

If you do not consider the host where you log in to be a secure host, do not use X forwarding. With X forwarding enabled, an attacker could authenticate via your `ssh` connection to intrude on your X server and sniff your keyboard input, for instance.

Caution

Buffer Overflows and Format String Bugs

As discussed in the section on local security, buffer overflows and format string bugs should be classified as issues concerning both local and network security. As with the local variants of such bugs, buffer overflows in network programs, when successfully exploited, are mostly used to obtain `root` permissions. Even if that is not the case, an attacker could use the bug to gain access to an unprivileged local account to exploit any other vulnerabilities which might exist on the system.

Buffer overflows and format string bugs exploitable over a network link are certainly the most frequent form of remote attacks in general. Exploits for these — programs to exploit these newly-found security holes — are often

posted on the security mailing lists. They can be used to target the vulnerability without knowing the details of the code. Over the years, experience has shown that the availability of exploit codes has contributed to more secure operating systems, obviously due to the fact that operating system makers were forced to fix the problems in their software. With free software, anyone has access to the source code (SuSE Linux comes with all available source codes) and anyone who finds a vulnerability and its exploit code can submit a patch to fix the corresponding bug.

DoS — Denial of Service

The purpose of this kind of attack is to force down a server program or even an entire system, something which could be achieved by various means: overloading the server, keeping it busy with garbage packets, or exploiting a remote buffer overflow.

Often a DoS attack is done with the sole purpose of making the service disappear. However, once a given service has become unavailable, communications could become vulnerable to so-called “man-in-the-middle attacks” (sniffing, TCP connection hijacking, spoofing) and DNS poisoning, explained below.

Man in the Middle: Sniffing, TCP Connection Hijacking, Spoofing

In general, any remote attack performed by an attacker who puts himself between the communicating hosts is called a “man-in-the-middle attack”. What almost all types of man-in-the-middle attacks have in common is that the victim is usually not aware that there is something happening. There are many possible variants, for example, the attacker could pick up a connection request and forward that to the target machine himself. Now the victim has unwittingly established a connection with the wrong host, because the other end is posing as the legitimate destination machine. The simplest form of a man-in-the-middle attack is called “sniffer” — the attacker is “just” listening to the network traffic passing by. As a more complex attack, the “man in the middle” could try to take over an already established connection (hijacking). To do so, the attacker would have to analyze the packets for some time to be able to predict the TCP sequence numbers belonging to the connection. When the attacker finally seizes the role of the target host, the victims will notice this, because they get an error message saying the connection was terminated due to a failure.

What often makes things easier for attackers is the fact that there are protocols which are not secured against hijacking through encryption, but only perform a simple authentication procedure upon establishing the connection.

Finally, we want to mention “spoofing”, an attack where packets are modified to contain counterfeit source data, mostly the IP address. Most active forms of attack rely on sending out such fake packets — something that, on a Linux machine, can only be done by the superuser (`root`).

Many of the attacks mentioned are carried out in combination with a DoS. If an attacker sees an opportunity to abruptly bring down a certain host, even if only for a short time, it will make it easier for him to push the active attack, because the host will not be able to interfere with the attack for some time.

DNS Poisoning

DNS poisoning means that the attacker corrupts the cache of a DNS server by replying to it with spoofed DNS reply packets, trying to get the server to send certain data to a victim who is requesting information from that server. To foist such false information onto the server in a credible way, normally the attacker must have received and analyzed some packets from it. Given that many servers are configured to maintain a trust relationship with other hosts, based on IP addresses or host names, such an attack may be successful in a relatively short time. On the other hand, it also requires quite an effort. In any case, the attacker will need a good understanding of the actual structure of the trust relationships between hosts. The attacker often needs to target a well-timed DoS attack at the name server, as well. Protect yourself by using encrypted connections that are able to verify the identity of the hosts to which to connect.

Worms

Worms are often confused with viruses, but there is a clear difference between the two. Unlike viruses, worms do not need to infect a host program to live. Rather, they are specialized to spread as quickly as possible on network structures. The worms that appeared in the past, such as Ramen, Lion, or Adore, make use of well-known security holes in server programs like `bind8` or `lprNG`. Protection against worms is relatively easy. Given that some time will elapse between the discovery of a security hole and the moment the worm hits your server, there is a good chance that an updated version of the affected program will be available on time. Of course, that is only useful if the administrator actually installs the security updates on the systems in question.

Some General Security Tips and Tricks

Information: To handle security competently, it is important to keep up with new developments and to stay informed about the latest security issues. One

very good way to protect your systems against problems of all kinds is to get and install the updated packages recommended by security announcements as quickly as possible. SuSE security announcements are published on a mailing list to which you can subscribe by following the link <http://www.suse.de/security>. The list suse-security-announce@suse.de is a first-hand source of information regarding updated packages and includes members of SuSE's security team among its active contributors.

The mailing list suse-security@suse.de is a good place to discuss any security issues of interest. Subscribe to it under the URL as given above for suse-security-announce@suse.de.

bugtraq@securityfocus.com is one of the best-known security mailing lists worldwide. We recommend that you read this list, which receives between 15 and 20 postings per day. More information can be found at <http://www.securityfocus.com>.

The following is a list of rules which you may find useful in dealing with basic security concerns:

- According to the rule of using the most restrictive set of permissions possible for every job, avoid doing your regular jobs as `root`. This reduces the risk of getting a cuckoo egg or a virus and protects you from your own mistakes.
- If possible, always try to use encrypted connections to work on a remote machine. Use “`ssh`” (secure shell) to replace `telnet`, `ftp`, `rsh` and `rlogin`.
- Avoid using authentication methods based on IP addresses alone.
- Try to keep the most important network-related packages up-to-date and subscribe to the corresponding mailing lists to receive announcements on new versions of such programs (`bind`, `sendmail`, `ssh`, etc.). The same should apply to software relevant to local security.
- Change the `/etc/permissions` file to optimize the permissions of files crucial to your system's security. If you remove the `setuid` bit from a program, it might well be that it cannot do its job anymore in the way it is supposed to. On the other hand, consider that, in most cases, the program will also have ceased to be a potential security risk. You might take a similar approach with world-writable directories and files.
- Disable any network services you do not absolutely require for your server to work properly. This will make your system safer, plus it prevents your users from getting used to a service that you had never

intended to be available in the first place (the legacy problem). Open ports, with the socket state LISTEN, can be found with the program `netstat`. As for the options, we suggest using `netstat -ap` or `netstat -anp`. The `-p` option allows you to see which process is occupying a port under which name.

Compare the `netstat` results with those of a thorough port scan done from outside your host. An excellent program for this job is `nmap`, which not only checks out the ports of your machine, but also draws some conclusions as to which services are waiting behind them. However, port scanning may be interpreted as an aggressive act, so do not do this on a host without the explicit approval of the administrator. Finally, remember that it is important not only to scan TCP ports, but also UDP ports (options `-sS` and `-sU`).

- To monitor the integrity of the files of your system in a reliable way, use the program `tripwire`. Encrypt the database created by `tripwire` to prevent someone from tampering with it. Furthermore, keep a backup of this database available outside your machine, stored on an external data medium not connected to it by a network link.
- Take proper care when installing any third-party software. There have been cases where a hacker had built a trojan horse into the tar archive of a security software package, which was fortunately discovered very quickly. Only install a binary package, if you have no doubts about the site from which you downloaded it.

SuSE's RPM packages are gpg-signed. The key used by SuSE for signing reads as follows:

```
ID:9C800ACA 2000-10-19 SuSE Package Signing Key <build@suse.de>
Key fingerprint = 79C1 79B2 E1C8 20C1 890F 9994 A84E DAE8 9C80
0ACA
```

The command `rpm -checksig package.rpm` shows whether the checksum and the signature of an uninstalled package are correct. Find the key on the first CD of the distribution and on most key servers worldwide.

- Check your backups of user and system files regularly. Consider that if you do not test whether the backup will work, it might actually be worthless.
- Check your log files. Whenever possible, write a small script to search for suspicious entries. Admittedly, this is not exactly a trivial task. In

the end, only you can know which entries are unusual and which are not.

- Use `tcp_wrapper` to restrict access to the individual services running on your machine, so you have explicit control over which IP addresses can connect to a service. For more information regarding `tcp_wrappers`, consult the man pages for `tcpd` and `hosts_access`.
- Design your security measures to be redundant: a message seen twice is much better than no message at all.

Using the Central Security Reporting Address

If you discover a security-related problem (please check the available update packages first), write an e-mail to security@suse.de. Please include a detailed description of the problem and the version number of the package concerned. SuSE will try to send a reply as soon as possible. You are encouraged to pgp encrypt your e-mail messages. SuSE's pgp key is as follows:

ID:3D25D3D9 1999-03-06 SuSE Security Team <security@suse.de>

Key fingerprint = 73 5F 2E 99 DF DB 94 C4 8F 5A A3 AE AF 22 F2 D5

This key is also available for download from: <http://www.suse.de/security>

YaST and SuSE Linux License Terms

YaST 2 Copyright (c) 1995 - 2001 SuSE GmbH, Nuernberg (Germany)

YaST 2 Copyright (c) 2002 SuSE Linux AG, Nuernberg (Germany)

The object of this licence is the Y_aST2 (Yet another Setup Tool 2) program, the name YaST, together with SuSE Linux, the Linux Distribution of SuSE Linux AG, all programme derived from Y_aST2 and all works or names derived in full or in part thereof together with the use, application, archiving, reproduction and passing on of Y_aST2, all programs derived from Y_aST2 and all works derived in full or in part thereof. The Y_aST2 program and all sources is the intellectual property of SuSE Linux AG within the meaning of the Copyright Law. The name YaST is a registered trademark of SuSE Linux AG. In the following SuSE Linux AG is the licensor and every user or processor of Y_aST2 or works derived in full or in part thereof, together with every person who reproduces, distributes or archives Y_aST2 or SuSE Linux, is the licensee of SuSE Linux AG.

The following licence terms are recognised as a result of the processing, use, application, archiving reproduction and dissemination of Y_aST2.

Only this licence gives the Licensee the right to use reproduce, to distribute or to amend Y_aST2 or works derived from it. These actions are forbidden by the copyright act, if this licence is not recognised. If this licence is recognised and complied with in full, it is also valid even without the written consent of the Licensee.

1. Usage

Y_aST2 and SuSE Linux may be used for personal and commercial purposes if the copyright and licence terms of the installed packages and

programmes are observed. The use of YqST2, even if a modified version is used, does *not* exempt in particular the Licensee from the duty to take due care with regard to the licence terms of the packages or programmes installed through YqST2 or works based on it.

2. Processing

All programmes derived from YqST2 and all works derived from it in full or parts thereof are to be filled on the opening screen with the clear information *Modified Version*. Moreover the operator give his name on the opening screen, stating that SuSE Linux AG is not providing any support for the *Modified Version* and is excluded from any liability whatsoever. Every amendment to the sources which are not conducted by SuSE Linux AG are deemed to be a *Modified Version*. The Licensee is entitled to change his copy from the sources of YqST2, whereby a work based on the YqST2 programme is created, provided that the following conditions are satisfied.

- (a) Every amendment must have a note in the source with date and operator. The amended sources must be made available for the user in accordance with section 3) together with the unamended licence.
- (b) The Licensee is obliged to make all work distributed by him which is derived as a whole or in part from YqST2 or parts of YqST2 to third parties as a whole under the terms of this licence without royalties.
- (c) The amendment of this licence by a Licensee, even in part, is forbidden.

SuSE Linux AG reserves the right to accept parts or all amendments of a modified version of YqST2 into the official version of YqST2 free of charge. The Licensee has no bearing on this.

3. Dissemination

It is forbidden to reproduce or distribute data carriers which have been reproduced without authorisation for payment without the prior written consent of SuSE Linux AG or SuSE Linux. Distribution of the YqST2 programme, its sources, whether amended or unamended in full or in part thereof, and the works derived thereof for a charge require the prior written consent of SuSE Linux AG.

All programmes derived from YqST2, and all works derived thereof as a whole or parts thereof may only be disseminated with the amended

sources and this licence in accordance with 2b. Making Yast2 or works derived thereof available free of charge together with SuSE Linux on FTP Servers and mailboxes is permitted if the licences on the software are observed.

4. Guarantee

No guarantee whatsoever is given for Yast2 or for works derived from it and SuSE Linux. The SuSE Linux AG guarantee only covers fault-free data carriers.

SuSE Linux AG will provide Yast2 and SuSE Linux *“as it is”* without any guarantee whatever that it is fit for a specific purpose or use. In particular SuSE is not liable for lost profit, savings not made, or damages from the claims lodged by third parties against the Licensee. SuSE Linux AG is not liable either for other direct or indirect consequential losses, in particular not for the loss or production of recorded data.

The observance of the respective licences and copyrights of the installed software is incumbent solely upon the user of Yast2 and SuSE Linux.

5. Rights

No other rights to Yast2 or to SuSE Linux are granted other than those negotiated in this licence. An infringement against this licence automatically terminates the rights of the Licensee. However the right of third parties who have received copies or rights under this licence from the Licensee, are not terminated as long as all parts of his licence are recognised and observed. If the Licensee is subject to conditions, or obligations as a result of a court judgement, patent terms, licence terms, or another reason, and these conditions or obligations contradict this licence as a whole or in part, the Licensee shall only be exempted in full or in part from this licence and its terms with the express prior written consent of SuSE. SuSE is entitled to withhold its consent without giving reasons.

6. Additional restrictions

If the distribution or use of Yast2 and SuSE Linux or parts of SuSE Linux is restricted in a state either by patents or by interfaces protected by copyright, SuSE Linux AG can specify an explicit geographic restriction of the distribution of Yast2 and SuSE Linux or parts of SuSE Linux, in which these states are fully or partially excluded from distribution. In such a case this licence includes the whole or partial restriction as if it was written down in this licence.

The GNU General Public License

GNU General Public License

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307, USA

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 675 Mass Ave, Cambridge, MA 02139, USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Foreword

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the *GNU General Public License* is intended to guarantee your freedom to share and change free software — to make sure the software is free for all its users. This *General Public License* applies to most of the *Free Software Foundation's* software and to any other program whose authors commit to using it. (Some other *Free Software Foundation* software is covered by the *GNU Library General Public License* instead.) You can apply it to your programs, too.

When we speak of “*free*” software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change

the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU General, Public License

Terms and Conditions for Copying, Distribution and Modification

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this *General Public License*. The "Program", below, refers to any such program or work, and a *work based on the Program* means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents

constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a

whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, “complete source code” means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or

binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the

only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The *Free Software Foundation* may publish revised and/or new versions of the *General Public License* from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the *Free Software Foundation*. If the Program does not specify a version number of this License, you may choose any version ever published by the *Free Software Foundation*.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the *Free Software Foundation*, write to the *Free Software Foundation*; we sometimes make exceptions for this.

Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

No Warranty

11. Because the program is licensed free of charge, there is no warranty for the program, to the extent permitted by applicable law. Except when otherwise stated in writing the copyright holders and/or other parties provide the program “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The entire risk as to the quality and performance of the program is with you. Should the program prove defective, you assume the cost of all necessary servicing, repair or correction.

12. In no event unless required by applicable law or agreed to in writing will any copyright holder, or any other party who may modify and/or redistribute the program as permitted above, be liable to you for damages, including any general, special, incidental or consequential damages arising out of the use or inability to use the program (including but not limited to loss of data or data being rendered inaccurate or losses sustained by you or third parties or a failure of the program to operate with any other programs), even if such holder or other party has been advised of the possibility of such damages.

End of Terms and Conditions

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

one line to give the program's name and a brief idea of what it does.

Copyright (C) 19yy *name of author*

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as

published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19yy name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for
details type 'show w'. This is free software, and you are welcome
to redistribute it under certain conditions; type 'show c' for de-
tails.
```

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items — whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
program 'Gnomovision' (which makes passes at compilers) writ-
ten by James Hacker.
```

```
Signed by Ty Coon, 1 April 1989
```

```
Ty Coon, President of Vice
```

This *General Public License* does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the *GNU Library General Public License* instead of this License.

Index

A

- ACLs
 - arranging 77
 - defining 75
- Adminhost 6, 10, 15–32
 - FAS 33
 - installing 16–22
 - troubleshooting 142
- administering
 - SSH 97
 - tools 15
- administrators
 - changing job of 4
- attacks 4
 - detecting 143–147
 - external 146
 - responding to 144

B

- BIND 123, 158–166
 - BIND8 160
 - BIND9 160
- booting
 - boot managers
 - YaST2 20
 - parameters 130

C

- CDs
 - Admin 5
 - burning 151
 - Live 5, 113–130
- certificates
 - CA 106
 - creating 107
 - deleting 107

- exporting 108
- importing 107
- managing 105
- chroot 125
- communication analysis 11
 - communication matrix 11
- compartment 125
- configuration files 126
 - .xsession 194
 - editing in FAS 109
 - firewalls.rc.config 187
 - host.conf 25
 - alert 25
 - multi 25
 - nospoof 25
 - order 25
 - trim 25
 - HOSTNAME 29
 - hosts 23, 126
 - hosts.allow 126
 - hosts.deny 126
 - inittab 126
 - ipsec.d 127
 - isdn.conf 126
 - lilo.conf 199
 - live-setup.conf 126
 - modules.boot 126
 - named 127
 - master 127
 - root.hint 127
 - slave 127
 - named.conf 127, 158–166
 - networks 24
 - nscd.conf 28
 - nsswitch.conf 25, 26

- ntp.conf 127
- pam.d 127
- permissions 206
- permissions.local 127, 200
- postfix 127
- ppp 126
- proxy-suite 127
- rc.config 23, 128, 173, 185, 200
- rc.config.d 128
- resolv.conf 28, 128, 172
- rinetd.conf 128
- route.conf 128
- runlevel.firewall 128, 129
- security 128
- shadow 128
- squid.conf 129, 173–179, 181
- squidguard.conf 180
- ssh 129
- sshd_config 194
- syslog.conf 129
- syslog.socks 129
- configuring
 - configuration disk 125
 - DNS 157
 - editing existing configurations .. 109
 - Squid 173
- connections
 - encrypting 121
- D**
- disks
 - configuration disk 125
- DNS 123
 - BIND 158–166
 - configuring 157–166
 - forwarder 61
 - forwarding 158
 - logging 161
 - options 160
 - reverse lookup 165
 - security and 205
 - starting 158
 - troubleshooting 158
 - zones 162
 - files 163
- documenting
 - configurations 36, 111
 - need for 16
- F**
- FAS 16, 33–111
 - base setup 41
 - DNS 61

- fwadmin 30, 34
- hard disks 42
- IP filters 53
- IP forwarding 53
- kernel 59
- logging 60
- NAT 53, 54
- network interfaces 43
- saving configurations 109
- starting 34
- users
 - creating 34
- feedback 155
- Firewall Administration System *see* FAS
- firewall host 10
 - booting 138
- firewalls 184, 186
 - configuring 33–111
 - documenting
 - need for 5
 - implementing 137–140
 - monitoring 111
 - new configurations 34
 - packet filters 184, 186
 - setups 12
 - SuSEfirewall 187
 - configuring 187
 - term origins 4
 - testing 138–140
- FReeS/WAN 121–123
 - security 122
- FTP
 - magic user 124
 - proxies 124
 - external 63
 - internal 66
- G**
- GNU
 - GPL 213–220
- H**
- hardware
 - required 114
- help 141
- HTTP
 - content filters 124
 - MIME 80
 - proxies 123–124
 - external 72
 - httpf 124
 - internal 74
 - tinyproxy 124

- proxy filters 78
- httpf 124

I

- ifconfig 142
- implementing
 - requirements 138
- installing
 - Adminhost 16–22
 - VPN 31
- IP addresses
 - masquerading 184
- IPsec 84, 121–123
- iptables 117–121
 - inserting rules 117

K

- kernel caps 125
- kernels
 - configuring 59

L

- Live CD 5, 113–130
 - advantages 147
 - reason for 114
 - services 115
 - troubleshooting 142
- log files 101
 - analyzing 99
 - IP filter 102
 - messages 158
 - Squid 173, 174
- log host 10
- logging 129
 - analyzing
 - interfaces 104
 - mail 104
 - analyzing log files 99
 - configuring 60
 - log host 10
 - value of 140

M

- mail 123
 - relaying 94
- maintaining 149–151
 - patches
 - getting 149
- masquerading 184
 - configuring with SuSEfirewall .. 187
 - ipchains 185
 - iptables 185
 - problems 186

- monitoring 111
 - need for 16

N

- name servers *see* DNS
- networks
 - attacks on 4
 - configuration files 23–30
 - configuring with YaST2 22
 - manual configuration 22–30
 - planning 10
 - troubleshooting 142–143
- NSS 26
 - databases 27

O

- OpenSSH *see* SSH

P

- package
 - bind8 166
 - howtoen 182
 - libcinfo 26
 - squidgrd 180
 - wget 150
- packet filters *see* firewalls
- partitioning
 - YaST2 18–20
- passwords
 - secure 34–35
- patches
 - getting 149
- Postfix 94, 123
- protocols
 - supported 6
- proxies
 - advantages 168
 - caches 168
 - FTP 124
 - external 63
 - internal 66
 - HTTP 123–124
 - external 72
 - httpf 124
 - internal 74
 - Squid 124
 - tinyproxy 124
 - Live CD and 115
 - parent 80
 - Squid 167
 - TCP 69
 - transparent 178

R

- root
 - password 20
- routing
 - masquerading 184
- RPM
 - security 207

S

- scripts
 - cipe 126
 - init.d 29, 129
 - inetd 29
 - network 29
 - nfsserver 30
 - portmap 29
 - route 29
 - sendmail 30
 - squid 172
 - ypbind 30
 - ypserv 30
 - SuSEconfig 23
- security 184, 195
 - attacks 204
 - booting 196, 199
 - bugs and 200, 203
 - DNS 205
 - file race 200
 - firewalls 184
 - local 197
 - networks 202
 - passwords 198
 - permissions 199
 - policy 11
 - protocols and 197
 - reporting problems 208
 - RPM signatures 207
 - serial terminals 196
 - server 6
 - Squid 168
 - SSH 190–195
 - tcpd 208
 - tips and tricks 205
 - viruses 201
 - worms 205
 - X and 202
- series
 - doc 26
 - n 180, 182
- Squid 124, 167–182
 - access controls 176
 - cache size 171
 - caches 168, 169

- Calamaris 181
 - configuring 173
 - CPU and 172
 - directories 172
 - features 168
 - hard disks and 170
 - log files 173, 174
 - object status 169
 - permissions 173, 176
 - protocols 168
 - RAM and 171
 - reports 181
 - SARG 182
 - security 168
 - SquidGuard 179
 - starting 172
 - transparent proxies 178
 - uninstalling 173
- SSH 125, 190–195
 - authentication mechanisms 193
 - configuring 97
 - daemon 192
 - key pairs 192, 193
 - scp 191
 - sftp 191
 - ssh 190
 - ssh-agent 193, 194
 - ssh-keygen 193
 - sshd 192
 - X and 194
 - support 141, 152–156
 - commercial 153
 - free 155–156
 - installation 152–153
 - services 153
 - support database 142
 - training 154
 - SuSEconfig 23
- ## T
- TCP
 - proxies 69
 - testing 139
 - configuration 110
 - external 140
 - internal 139
 - time
 - synchronizing 99
 - tinyproxy 124
- ## U
- updating
 - patches 149

upgrading
- VPN 31

users
- fwadmin 34

V

VPN 31, 84
- configuration 84
- masquerading 88
- NAT 88

W

web service
- proxies 123–124

X

X
- security 202
- SSH and 194

xntpd
- configuring 99

Y

YaST2
- boot manager 20
- graphics card 21
- hard disks 18
- installing 21
- keyboard 17
- language 16
- monitor 21
- mouse 17
- navigating 16
- network 22
- partitioning 18–20
- root password 20
- time zone 17