


Secure Your E-mail Server on IBM @server i5 with Linux

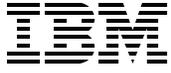
Understanding security issues for
network and e-mail server

Linux open source solutions to
secure your e-mail server

Linux-based ISV solutions to
secure your e-mail server



Yessong Johng
Alex Robar
Colin McNaught
Senthil Kumar



International Technical Support Organization

Secure Your E-mail Server on IBM @server i5 with Linux

October 2005

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Archived

First Edition (October 2005)

This edition applies to IBM i5/OS V5R3, SUSE LINUX Enterprise Server 9, and Red Hat Enterprise Linux AS Version 4.

© Copyright International Business Machines Corporation 2005. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this Redpaper	ix
Become a published author	xi
Comments welcome	xi
Part 1. Open Source Solutions for Network Security	1
Chapter 1. Understanding and planning e-mail server security	3
1.1 Concepts: Securing e-mail servers	4
1.1.1 Linux-based firewall	6
1.1.2 E-mail security	7
1.2 Scenarios: Securing e-mail server	8
1.2.1 Open source protection and open source mail delivery	8
1.2.2 Open source protection and Domino	9
1.2.3 ISV protection, open source filtering, and open source mail delivery	10
1.3 Planning: Securing e-mail server	11
1.3.1 OSS versus ISV solutions for network security mechanisms	11
1.3.2 Direct I/O for firewall	11
1.3.3 Choice of e-mail server	12
1.3.4 Support contracts	12
1.3.5 De-militarized zone	13
1.3.6 Planning worksheet	13
1.4 Types of attacks and protection mechanisms	15
Chapter 2. Linux installation	17
2.1 Linux installation overview	18
2.1.1 Required or helpful tools	18
2.1.2 Installation notes	18
2.2 Setting up the partitions	20
2.2.1 Creating a logical partition using the HMC	20
2.2.2 Set up the i5/OS partition virtual I/O	40
2.2.3 Working with network servers	44
2.3 Installing Linux	50
2.3.1 Installing SLES9	50
2.3.2 Installing RHEL4	80
Chapter 3. Locking down the Linux firewall partition	107
3.1 Hardening Linux	108
3.1.1 Bastille Linux	108
3.1.2 Removing unnecessary servers	156
3.1.3 Altering insecure defaults	156
3.2 iptables rules	159
3.2.1 Understanding iptables	160
3.2.2 Initial iptables setup	162
3.3 grsecurity kernel patch	166
3.4 Security-Enhanced Linux (SELinux)	171

3.5 Snort	171
3.5.1 Installing libpcap 0.9.0-096	171
3.5.2 Installing Perl Compatible Regular Expressions (PCRE) 5.0	172
3.5.3 Installing Snort 2.3.1	172
3.5.4 Configuring Snort	173
3.6 Rootkit hunter	175
Chapter 4. E-mail Security tools installation and configuration	177
4.1 Postfix	178
4.1.1 Preparing to install Postfix	178
4.1.2 Updating Postfix	181
4.1.3 Postfix configuration files	183
4.1.4 Configuring Postfix	184
4.2 qmail	185
4.2.1 Overview of qmail installation	186
4.2.2 Preparing to install qmail	186
4.2.3 Installing qmail	188
4.2.4 Configuring qmail	195
4.3 Clam AntiVirus	202
4.3.1 Installing Clam Antivirus	203
4.3.2 Configuring Clam AntiVirus for Postfix	208
4.3.3 Configuring Clam Antivirus for qmail	211
4.3.4 Adding Clam to system boot	213
4.4 SpamAssassin	213
4.4.1 Installing SpamAssassin	214
4.4.2 Overview: Configuration of SpamAssassin	219
4.4.3 Configuring SpamAssassin for Postfix	222
4.4.4 Configuring SpamAssassin for qmail	223
4.4.5 Adding SpamAssassin process to system startup	224
4.4.6 Installation and configuration of qmail-scanner	224
Part 2. ISV solutions	229
Chapter 5. StoneGate Firewall and VPN for iSeries	231
5.1 Why commercial firewalls	232
5.1.1 Open source verses commercial firewalls	232
5.1.2 Internal verses external firewalls	232
5.2 The StoneGate Firewall	233
5.2.1 Multi-Link Technology	234
5.2.2 Virtual Private Networking (VPN)	234
5.2.3 Firewall configuration and maintenance	236
5.2.4 StoneGate management client	236
5.2.5 Centralized management center: SMC	237
5.2.6 Reporting tools	238
5.2.7 Clustering and VPN high availability	239
5.3 Planning and preparation for the StoneGate Firewall	241
5.3.1 Hardware and software requirements for eServer i5 and SMC	242
5.4 Other scenarios	243
5.4.1 Integration environment	243
5.4.2 External Firewall Cluster and integration environment	244
Chapter 6. MPP Service Provider Edition: Foundation of E-mail Filtering Services	247
6.1 The core of your e-mail service: Policy-based service offerings	248
6.2 Selected MPP-based service offerings	248

6.3 MPP engine and e-mail server support	248
6.4 Managing MPP	249
6.5 Conclusion	249
Chapter 7. Bytware's StandGuard Anti-Virus for Linux	251
7.1 The McAfee virus scanning engine	252
7.2 Additional StandGuard Anti-Virus features	252
7.2.1 Command-line interface and graphical user interface.	252
7.2.2 Automatic download of virus definitions (.DAT files).	252
7.2.3 Automatic download of software updates and fixes	253
7.2.4 Built-in scheduling features for scanning and updating.	253
7.2.5 Network enabled	254
7.2.6 Extensive logging capabilities	255
Appendix A. Tips and techniques	257
SpamAssassin: Language options.	258
ok_language	258
ok_locales	258
gmail control files	259
Packages: Links to download.	260
Related publications	263
IBM Redbooks	263
Online resources	263
How to get IBM Redbooks	264
Help from IBM	264

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX 5L™
AIX®
AS/400®
Domino®
IBM®
ibm.com®
IBM @server

IBM @server
i5/OS™
iSeries™
Lotus Notes®
Lotus®
Notes®
OS/400®

POWER™
POWER5™
Redbooks™
Redbooks (logo) ™
WebSphere®

The following terms are trademarks of other companies:

Power Management, Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Outlook, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Pentium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.


Preface

This IBM Redpaper will help you design a solution to protect your e-mail servers on IBM® eServer™ i5 using various security solutions of Linux®. Like many of solutions of Linux, security implementation can be done either using free software solutions (that is, Open Source Software (OSS) based solutions) or commerce software solutions (that is, Independent Software Vendor (ISV) based solutions).

This redpaper has two main parts: Part 1, “Open Source Solutions for Network Security” on page 1, covers OSS-based solutions and Part 2, “ISV solutions” on page 229, covers ISV-based solutions. In either case, the goal is how to protect your e-mail servers from various attacks.

The team that wrote this Redpaper

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



Yessong Johng is an IBM Certified IT Specialist at the IBM International Technical Support Organization, Rochester Center. He started his IT career at IBM as a S/38 Systems Engineer in 1982 and has been with S/38, AS/400®, and now iSeries™ for 20 years. He writes extensively and develops and teaches IBM classes worldwide on the areas of IT Optimization whose topics include Linux, AIX®, and Windows® implementations on iSeries. He is also interested in the e-business area, especially with WebSphere® implementations on iSeries.




Alex Robar is an IT Specialist from ASTECH Solutions, Inc., an IBM Business Partner based in Aurora, Ontario, Canada. He has three years of experience in the field of installing and administering Linux on iSeries LPARs. His skills include training material development, Linux installation, administration and security, iSeries LPARs, and Web development. His expertise has been used to provide solutions for many client problems.



Colin McNaught is a Senior Consultant, Imtech ICT Brocom. Colin has provided consultancy and implementations covering all aspects of ICT infrastructure for the IBM @server for over fifteen years in the UK and Benelux. He is a specialist in the areas of security, covering firewall on iSeries to Digital Certificates and SSL, for availability from Backup/Recovery to High Availability, and for consolidation of Windows, Linux and AIX, in iSeries and i5 environments, as well as the optimization of these environments. He also provides educational classes in all of the above to clients. Colin is an IBM Certified iSeries Systems Expert for Windows, Linux integration, and Domino®, as well as a Certified StoneGate Architect for the firewall from

StoneSoft.



Senthil Kumar is a Software Engineer at IBM Software Labs, India. He has over five years of experience in UNIX® system administration on Red Hat and SUSE Linux on various hardware architectures. His areas of expertise include server security, mail servers, DNS servers, database servers, Web servers, and file and network servers. He holds a Bachelor's Degree in Mechanical Engineering from the University of Madras and is a Red Hat certified professional in RHEL.

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

Julie Czubik
International Technical Support Organization, Poughkeepsie Center

Fant Steele
Michael Ranweiler
IBM Rochester

Edmund Stanton
IBM Cambridge

Robert Macfarlane
RJIS

Christopher Watts
Linian Inc.

Hanna Lassila-Sramek
Stonesoft Corporation

Michael Katz
RAE Internet

Lennie Broich
Chris Grant
Bytware, Inc.

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and clients.

Your efforts will help increase product acceptance and client satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Obtain more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829

Archived




Part 1

Open Source Solutions for Network Security

This part discusses the Open Source Software (OSS) based network security mechanism.

Archived



Understanding and planning e-mail server security

This chapter discusses the ideas and concepts covered in this redpaper. The scenarios in which this information would be useful are also covered.

1.1 Concepts: Securing e-mail servers

This redpaper uses the power of the IBM @server i5 to protect your organization by securing your network against all sources of threats, and filtering your mail to eliminate viruses and spam.

The IBM i5/OS™ and the open source Linux operating system are inherently secure. As such, they are the perfect choice to secure your infrastructure. The i5/OS partition provides the ability to assign virtual hardware, such as virtual Ethernet adapters, to multiple Linux operating systems running in Logical Partitions (LPARs). This allows for secure communication between servers, and eliminates the possibility of interception or alteration of critical data once it has reached your network.

The Linux operating system provides the ability to allow as much or as little freedom to accounts on the system as is desired. This allows for one Linux system to function as a firewall, disallowing all other processes from running, while another Linux system scans, sorts, filters, and delivers mail to hundreds of different users, unencumbered by the process restrictions of the firewall.

This redpaper aims to combine these features to provide you with first class security using a dual Linux LPAR setup. This setup will protect your network, and filter incoming mail messages that contain spam content or viruses.

Figure 1-1 on page 5 is a visual representation of the security measures suggested in this redpaper.




Figure 1-1 Suggested security measures

As a first line of defense, all Ethernet interfaces (physical or virtual) will be monitored by the Snort network Intrusion Detection System (IDS), running on a Linux LPAR (LINUXFW). Any suspicious activity will be logged for further scrutiny.

Running on the same partition as the Snort IDS will be the built-in Linux netfilter/iptables firewall system. The iptables system will use a customized set of rules to allow only authorized packets through. Authorized packets will be forwarded to their correct destination once they are allowed through the firewall. Unauthorized packets will simply be dropped.

On a second Linux LPAR, LINUXST, the Postfix mail transfer agent (MTA), Clam AntiVirus (ClamAV), and SpamAssassin software packages will be installed. Mail received by the firewall from your internal LAN or your Internet connection will be forwarded to the LINUXST partition. Postfix will then run each message through the ClamAV and SpamAssassin filters. If a message is identified as containing a virus, the message is quarantined immediately. If a message is identified as spam, the subject line will be tagged, and the message will then be sent back to Postfix. Depending on your system setup, Postfix will then either deliver the message to a local user mailbox (if you are using a Linux e-mail server such as Postfix), or redirect the message back to the LINUXFW partition, which will then route the message to any e-mail server running on the i5/OS partition, the Domino server, for example.

Note: You might wonder why the interface to LAN, eth1, is also connected through the LINUXFW partition instead of directly connected to the i5/OS partition. LAN is our internal network and by definition it is a trusted network, correct? Unfortunately, the answer is not really. More security violation is committed by internal people than external people. Therefore, it is a very sound practice to treat your internal network as yet another untrusted network.

For additional security, all the accesses to the i5/OS partition will be routed through the firewall partition. As such, there will be no unsecured Ethernet interfaces entering the eServer i5 system. The result of all these security measures is a rock solid setup with the capability to perform packet and e-mail filtering for thousands of users on your network.

1.1.1 Linux-based firewall

The setup detailed in this redpaper uses the firewall capabilities of Linux to scan all packets that it receives from your internal LAN or your connection to the Internet. Any dangerous or malicious packets will be dropped. All access to any service on the Internet from users on your LAN will go through the firewall, which will protect your network against malicious external users. Any attack that comes from your internal network will also be blocked.

Access to the i5/OS partition will also be restricted through the firewall. This will disallow spoofed addresses from accessing your i5/OS. As an additional benefit, all accesses will be logged. In the event of an attack on any services running on the i5/OS, you will have the ability to determine the security hole an attacker was trying to exploit.

To ensure that nobody can access the Linux firewall who should not be able to, the operating system will be hardened, and patched with the Grsecurity and Security Enhanced Linux (SELinux) kernel-level patches. These patches will enforce strict security policies that make unauthorized access to your system very difficult. All unnecessary services will be disabled, and only secure communication protocols will be used for administration. A network intrusion detection system will also be utilized to monitor all network traffics. Any suspicious activity will be logged. A visual representation of the path of a packet through the Linux-based firewall would look like Figure 1-2 on page 7.




Figure 1-2 Packet path through Linux security system

1.1.2 E-mail security

The efficiency of the Linux operating system combined with other Open Source Software (OSS) packages creates a mail filtering system that is capable of handling large volumes of mail efficiently and accurately. OSS mail transfer agents will facilitate the transfer of mail through a constantly updated virus scanner and a spam filtering system that uses a wide variety of techniques to identify spam. This system will minimize the number of false-positives identified as carrying a virus or spam content. The isolation of the LPAR containing the Linux e-mail security tools also ensures that tampering with the filtering or scanning rules is a tremendously difficult task. A visual representation of an e-mail traversing the implemented OSS scanning and filtering agents would look like Figure 1-3 on page 8.




Figure 1-3 E-mail path through e-mail security system

1.2 Scenarios: Securing e-mail server

The information contained in this redpaper can be applied to many different infrastructure scenarios. It is possible to use an entirely OSS solution for your e-mail. You may prefer the software of an Independent Solution Provider (ISV). This section provides example scenarios that could be used to secure your infrastructure.

1.2.1 Open source protection and open source mail delivery

This redpaper can provide a good base for setting up a Linux-based mail server behind a secure firewall that protects the rest of your infrastructure. In this scenario, mail users would be set up on the e-mail security tools partition. Mail would be accepted by the firewall partition and forwarded via virtual Ethernet to the e-mail security tools partition. The mail would be scanned for viruses, and processed for signs that it is spam. Passing both these checks, the mail is delivered to the appropriate local mailbox on the system. Any user wishing to retrieve their mail would send their request to the firewall partition, which would forward the connection over virtual Ethernet to the e-mail security tools partition. The open source mail transfer agent on the e-mail security tools partition would then provide the user with their mail. The blue line in Figure 1-4 on page 9 illustrates the path of a mail message traversing the OSS protection, OSS delivery scenario.




Figure 1-4 Mail path through OSS protection, OSS delivery scenario

1.2.2 Open source protection and Domino

If you have an ISV mail server such as Domino already running on your network, this redpaper provides the perfect way to filter your messages in addition to blocking unauthorized access to any other network services. In this scenario, mail users are defined via the setup of the Domino software. Mail is accepted by the firewall partition and forwarded to the e-mail security tools partition over virtual Ethernet. The mail is then scanned for viruses, and processed for signs that it is spam. Passing both these checks, the mail is forwarded back to the firewall partition over the same virtual Ethernet network. The firewall partition will then redirect the mail up to the Domino server over a different virtual Ethernet network, which will deliver the mail to the proper user mailbox. Any user wishing to retrieve their mail would send their request to the firewall partition, which would forward the connection over the virtual Ethernet to the i5/OS partition, where Domino would provide the user with their mail. The blue line in Figure 1-5 on page 10 illustrates the path of a mail message traversing the OSS protection, Domino delivery scenario.

Note: Native solutions for virus scanning and spam filtering do exist for Domino servers running on the i5/OS: They are Symantec AntiVirus/Filtering for Domino 3.0 for OS/400® and Trend Micro ScanMail for Lotus® Notes® 2.6 for OS/400. For details, refer to the IBM Redbook, *Domino 6 for iSeries Best Practices Guide*, SG24-6937.




Figure 1-5 Mail path through OSS protection, Domino delivery system

1.2.3 ISV protection, open source filtering, and open source mail delivery

This redpaper also provides information for securing your infrastructure using the StoneGate firewall and the VPN system. In this scenario, mail users would be set up on the e-mail security tools partition. Incoming mail would be accepted by the StoneGate firewall, which is installed in the Linux LPAR partition, and forwarded via the virtual Ethernet to the e-mail security tools partition. The mail would be scanned for viruses, and processed for signs that it is spam. Passing both these checks, the mail would be delivered to the appropriate local mailbox on the system. The StoneGate firewall would forward the connection over the virtual Ethernet to the e-mail security tools partition. The open source mail transfer agent on the e-mail security tools partition would then provide the user with their mail. The blue line in Figure 1-6 on page 11 illustrates the path of a mail message traversing the ISV protection, OSS filtering, OSS delivery scenario.




Figure 1-6 Mail path through ISV protection, OSS filtering, OSS delivery scenario

1.3 Planning: Securing e-mail server

This section provides some considerations or decision making points when it comes to the implementation and security of an e-mail server.

1.3.1 OSS versus ISV solutions for network security mechanisms

Like most software applications, there are many choices for the e-mail servers and filtering agents that perform the tasks laid out in this redpaper. The first step in choosing which software applications best suit your infrastructure is to decide whether you should use software released under an open source license, or whether you should use software provided by an ISV. There are advantages to both setups.

OSS has the potential to save you substantial sums of money, as all OSS is released for free. There is also the added benefit that the source code for the software is available for free download, allowing you to tailor the software to your needs if you would like functionality that is not present by default.

ISV-provided software has the advantage of proven testing procedures and operations. An ISV will support their product when purchased from them, making software-related problems much easier to fix.

1.3.2 Direct I/O for firewall

The LPAR setup in this redpaper details using virtual I/O to run your Linux operating systems. If you are using your firewall as a gateway for your entire network, you may choose to use direct I/O instead of virtual I/O. If direct I/O is used, in the event that the i5/OS needs to be taken offline, the firewall will remain functional. If virtual I/O is used, the firewall LPAR will shut down with the i5/OS.

1.3.3 Choice of e-mail server

Depending on the e-mail server you choose, your security plan to secure your e-mail server can change. If you have an existing enterprise mail server, you may choose to keep your existing infrastructure and simply secure it. If you are setting up an infrastructure that does not have an existing mail server, or you would like to simplify your setup, you may choose to use an entirely open source delivery option.

Domino

Domino is the IBM enterprise e-mail server offering. Domino provides messaging, calendar, and scheduling capabilities, all running natively on the i5/OS.

On the one hand, the security mechanisms provided in this redpaper work well with Domino, provided your strategy is to consolidate your network security into Linux LPARs. On the other hand, you may opt to run all e-mail server and security mechanisms on the i5/OS. In this case, you could use either Symantec AntiVirus or Trend Micro AntiVirus solutions for mail protection, both of which run natively with the Domino server on the i5/OS.

For further information for these ISV solutions, refer to *Domino 6 for iSeries Best Practices Guide*, SG24-6937.

OSS solutions

There are many OSS mail transfer agents that enable the construction of a multi-user mail server infrastructure for little or no cost. Sendmail, while a reliable choice, is an older MTA that tends to be difficult to configure. Postfix is an alternative for Sendmail that aims for 100 percent compatibility for all Sendmail directives and hooks, while maintaining an easy configuration setup, and faster execution. The qmail package provides the same functionality as Sendmail and Postfix, while using a different configuration scheme entirely. This redpaper provides instructions for the setup and configuration of Postfix and qmail.

Bynari

Bynari's Insight Family of products provides an integrated e-mail server for the small and medium businesses.

The Insight Server runs under Linux with i5/OS in a partition and provides an integrated e-mail environment with tools like backup and recovery, migration wizards for Exchange and other IMAP servers, anti-virus, and anti-spam applications.

The Insight connector supports groupware features and functions with Outlook® 98, 2000, 2002, and 2003. Share calendars, contacts, public folders, tasks, and other features/functions. The Web client supports collaborative features that allow integration with Outlook calendar, contacts, tasks, and public folders.

To get a complete list of supported features and functions, go to:

<http://www.bynari.net/>

1.3.4 Support contracts

Clients have a choice of where they get support for Linux.

Users can get the support from the distributors.

- For SLES9 users, refer to the following link for information about the support from Novell:

http://support.novell.com/linux/linux_server_support.html

Users can order this support from IBM via econfig with their Linux distribution or order it from Novell.

- ▶ For RHEL4 users, refer to the following link for information about the support from Red Hat:

<http://www.redhat.com/software/rhel/compare/server/>

Users can order this support from IBM via econfig with their Linux distribution or order it from Red Hat.

- ▶ Or, users can get the support from IBM Global Services (IGS) via Support Line:

<http://www.ibm.com/services/us/index.wss/offering/its/a1000030>

Users cannot order Support Line for Linux in econfig.

Note: The i5/OS Support Line provides support for the integration of Linux on iSeries servers, for example, how to create a partition and how to use virtual I/O. This Support Line does not provide support for Linux itself.

1.3.5 De-militarized zone

A De-Militarized Zone (DMZ) is where servers usually reside on your network. Network traffic routed to devices on the DMZ completely bypasses the firewall. Usually, the DMZ has some form of firewall protecting any devices within it. If your network has a DMZ, or you would like to add one to your existing infrastructure, simply add a third physical Ethernet adapter to your firewall partition. In addition, the firewall rules must be altered so that all packets that do not match any specifically set rules are accepted and forwarded to the IP address of the DMZ.

1.3.6 Planning worksheet

Table 1-1 is a planning worksheet. It is a good idea to fill this out now, as it will be referenced many times during installation and setup. To use the planning worksheet, simply look in the Item column to find the row of the particular component you require. Once you have found the appropriate row, find the value of the cell in that row that is in the column of the partition you are currently setting up. For example, if you are setting up the firewall partition, and you require the Network server text description, you would look in cell E2. The value is *Linux firewall network server description*. Note that cells with a value of N/A do not need to be filled out, and are not applicable to the current situation.

Attention: The planning worksheet is a good checkpoint to make sure you are ready for the installation process. If there are any blank cells for which you do not know the value, you should discover their value prior to starting the installation. Failure to do so may result in an incomplete Linux operating system installation or invalid settings.

Table 1-1 Planning worksheet

	Item	1	2	3
A	Partition use	i5/OS	Firewall	Security
B	LPAR Name	i5/OS	LINUXFW	LINUXST
C	Profile Name		Default	Default
D	Network server description name (NWSD)	N/A	LINFWSD	LINSTSD

E	Network server text 'description' (NWS D)	N/A	Linux firewall NWS D	Linux security tools NWS D
F	Network storage space name (NWSSTG)	N/A	LINUXFW	LINUXST
G	Network storage space text 'description' (NWSSTG)	N/A	Linux firewall storage space	Linux security tools storage space
H	Network storage space size	N/A	6GB	8GB
I	Virtual console IP	N/A		
J	Linux 'root' password	N/A		
K	Gateway	N/A		N/A
L	Primary DNS			
M	Secondary DNS			
N	Host Name		LINUXFW	LINUXST
O	Domain	N/A		
P	Regular username	N/A		
Q	Regular user password	N/A		
	Firewall to Security Tools Interface (Virtual Ethernet)			
AA	Network device name	N/A	eth2	eth0
AB	IP Address	N/A	10.1.1.10	10.1.1.20
AC	Subnet Mask	N/A	255.255.255.128	255.255.255.128
	Firewall LAN Interface (Physical Ethernet)			
BA	Network Device Name	N/A	eth1	N/A
BB	IP Address	N/A		N/A
BC	Subnet Mask	N/A		N/A
	Firewall External Interface (Physical Ethernet)			
CA	Network device name	N/A	eth0	N/A
CB	IP Address	N/A		N/A
CC	Subnet Mask	N/A		N/A
	Firewall to i5/OS Interface (Virtual Ethernet)			
DA	Network device name		eth3	N/A
DB	IP Address	172.27.72.20	172.27.72.10	N/A
DC	Subnet Mask	255.255.255.128	255.255.255.128	N/A

1.4 Types of attacks and protection mechanisms

There are many different types of attacks, and consequently many different ways to defend against these attacks. Table 1-2 provides a reference for which tools are mentioned in this redpaper, and which attacks they will help defend against.

The left column lists potential attacks that could be launched against your system. The top row lists products. An X in any given cell means that the product listed at the top of that column protects against the attack listed in the left-most cell in that row.

Table 1-2 *Types of attacks and protection mechanisms*

	SpamAssassin	ClamAV	Linux Netfilter Firewall	Hardened Linux Installation
Spam	X			
Hoax E-mail	X			
Virus Attached to E-mail		X		
Internet Worm			X	
Denial of Service				X
Distributed Denial of Service			X	
Unauthorized Network Usage			X	X
Rootkit				X
Scanners			X	
Trojan		X	X	

Archived



Linux installation

This chapter discusses setting up multiple Linux operating systems on the IBM @server i5.

2.1 Linux installation overview

Setting up multiple Linux operating systems to run concurrently is a process that has been greatly simplified on the eServer i5. This chapter discusses the initial LPAR setup on the eServer i5 and then installation of the Linux operating systems.

2.1.1 Required or helpful tools

The setup and installation process for the secure e-mail server that will be built using this redpaper requires specific tools. These tools allow creation and maintenance of LPARs in addition to providing remote access to the Linux servers via the telnet and SSH protocols.

PuTTY

Putty is an open source implementation of the telnet and SSH protocols along with an xterm window. PuTTY runs on Windows and Unix platforms. We recommend the usage of PuTTY for all installations, as it supports the full implementation of the telnet and SSH protocols, in addition to being able to accurately display Linux console output.

PuTTY does not require installation. It is a single executable file that can be downloaded from:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

Hardware Management Console (HMC)

The instructions in this redpaper require the use of an HMC. The HMC allows configuration and management of server partitions and capacity on demand for POWER5™ servers. The HMC allows for a wizard-guided LPAR profile setup processes.

Important: i5/OS has been enhanced with support for virtual partition management to enable the creation and management of Linux partitions without the requirement for a Hardware Management Console (HMC). With the Virtual Partition Manager (VPM), an eServer i5 server can support one i5/OS partition and up to four Linux partitions.

Nevertheless, we cannot use VPM for setting up firewalls because the Linux partitions must use virtual I/O resources that are owned by the i5/OS partition.

2.1.2 Installation notes

Installing Linux on an LPAR on an eServer i5 is a different process than a standard installation of Linux on other system architectures. These notes will help you understand and navigate through the installation process.

Text-based administration

Installation of the Linux operating system, as well as all related software used in this redpaper, is done entirely via a text interface. Navigate through the text-based screens using the following key combinations:

- ▶ Ctrl+C is used to abort screens. We recommend that this not be used unless it is absolutely necessary, as it can leave an installation half finished.
- ▶ Ctrl+L is used to redraw the screen, if there appear to be rendering errors.
- ▶ Ctrl+H is used in place of the Backspace key.
- ▶ Tab is used to move between elements in an installation program.

If your system seems to have difficulties sending the Tab key press to the server, you may wish to use Ctrl+F instead.

- ▶ Spacebar is used as a toggle to select or deselect check boxes or radio buttons.
- ▶ Enter accepts a selection or activates a highlighted button.
- ▶ The `clear` command can be entered at the Linux shell to remove all previous commands and output from the screen.

In Linux scripts and configuration files, when a pound sign (#) precedes a line of text, it means that the text following the symbol is a comment. A comment is human-readable text that is meant to highlight or explain a certain part of a file for anyone who might be reading it. Comments are ignored by the system when the file is being read or executed. If you are instructed to *comment out* a line of text, you should place a pound sign (#) at the beginning of the appropriate line. If you are instructed to *uncomment* a line of text, simply remove the pound sign (#) from the beginning of the line.

Specific versions

The instructions in this redpaper list specific versions of all operating systems and software packages that are used. We recommend that you use the specified versions of all software. Alternate versions may have different installation or configuration procedures, and may not function as intended.

Distributions

This redpaper supports the usage of both SUSE Linux Enterprise Server 9 and RedHat Enterprise Linux 4. The installation procedure for each distribution is different. There are separate sets of instructions for installing SUSE and installing RedHat. After the installation of the operating system, software installation on the two distributions is nearly identical, so there is only one set of instructions. When there is a discrepancy between the way a task is accomplished, there is a note that informs you of how that task will be accomplished on both SUSE and RedHat.

Basic security

The Linux operating system can provide your network with world class security. That security is worthless, however, if decent security practices are not followed by all users of your servers. We recommend that you follow the general good security practices listed here:

- ▶ Passwords should be changed often. The instructions in this redpaper do not enforce password aging due to the inherent problems with this practice. However, we recommend that all accounts on the system have their password changed weekly. Passwords should follow these guidelines:
 - Use a combination of both letters and numbers.
 - Do not use common names, or names of those close to you (such as a spouse, child, or pet).
 - Do not use phone numbers, social security numbers, or birth dates. Not only is this easily crackable if an attacker knows you personally, but it is also a danger to your identity if your account is cracked, and an attacker can read your password.
 - Do not use the same name as your login.
 - Do not use any words that can be found in the dictionary, either local or foreign.
 - *Never* use a blank password. All the security in the world is absolutely worthless if you leave any password on your system blank.
- ▶ Log out when you have finished working with your system or are leaving your terminal. An open console is an open invitation to damage, even if the damage is accidental. The

virtual console system used to setup the Linux operating system initially is especially dangerous, as it is a shared console. Multiple users can connect to the console and watch as commands are entered and output is displayed. Watch console output, ensure that you know who is connected, and log out when you are finished.

2.2 Setting up the partitions

Partition setup is a three-step process:

1. Create an LPAR using the HMC.
2. Set virtual input/output (I/O) settings on the i5/OS partition.
3. Create and link a Network Server Storage Space (NWSSTG) and a Network Server Description (NWS D).

This section has three subsections that describe each step in further detail.

2.2.1 Creating a logical partition using the HMC

Partition profiles are a new concept introduced with the eServer i5 LPAR. If you are new to LPAR, you will need to be familiar with this new concept. Refer to *Logical Partitions on IBM PowerPC: A guide to working with LPAR on Power5 IBM eServer i5 servers, SG24-8000*.

The HMC has a wizard for both partition and partition profile creation. The following steps describe how to use this wizard. Note that to create a partition or partition profile, you must have super administrator or operator privileges on the HMC.

Important: The following steps must be completed twice for both LPARs to be created. On the first iteration of the steps, use the values in column 2 on the planning worksheet to create the firewall partition. On the second iteration of the steps, use the values from column 3 to create the security tools partition.

1. In the Navigation Area of your HMC, select the name of your server, then select **Server and Partition** → **Server Management**.
2. In the Server and Partition: Server Management pane, expand your server name. Right-click **Partitions**, and select **Create** → **Logical Partition**, as shown in Figure 2-1 on page 21.




Figure 2-1 Creating an LPAR

3. The Create Logical Partition Wizard appears.
 - a. Leave the Partition ID value as the default.
 - b. Fill in the Partition name from row B from the planning worksheet.
 - c. Click **AIX or Linux** under the Partition environment heading, as shown in Figure 2-2 on page 22. Click the **Next** button.




Figure 2-2 LPAR - Creating a new environment

- d. The Workload Management Groups options appear. Select **No**, as shown in Figure 2-3 on page 23. Click **Next** to continue.




Figure 2-3 LPAR - Workload Management Groups

- e. Create an LPAR profile. The profile object specifies the characteristics of the partition, such as allocated memory, processors, I/O devices, and slots.

As shown in Figure 2-4 on page 24, enter the Profile name from row C on the planning worksheet. Click **Next** to continue.

Important: Do not check the “Use all the resources in the system” option. If this box is checked, the partition will try to access and utilize all the physical resources in the system when it is activated.



Figure 2-4 LPAR - Create Logical Partition Profile

- f. Specify memory options. Memory options are very environment specific. If there will be a lot of traffic traversing your network, it is best to increase these settings. For smaller, low traffic networks, lower settings will suit your needs. As a general rule, we recommend setting the Minimum memory to 512 MB.

We use 512 MB as our Minimum memory, 1 GB as our Desired memory, and 3 GB as our Maximum memory, as shown in Figure 2-5 on page 25.

Once your memory settings have been entered, click **Next** to continue.




Figure 2-5 LPAR - Memory settings

4. Choose whether to use shared processing units or a dedicated processor for the LPAR you are creating. Again, this option is very much based upon your network. If there will be a very high volume of traffic, you may choose to share a processor solely between the firewall and the security tools partition. We recommend that you select **Shared**, as shown in Figure 2-6 on page 26. The shared option will work well for most setups. It would be very rare to have an entire processor dedicated to a single firewall or security tools partition.

Once you have made your selection, click the **Next** button to continue.




Figure 2-6 LPAR - Selecting a dedicated or shared processor

- g. If you selected the shared processing option, you will be presented with the Processing Settings window. Here you enter the Desired processing units, the Minimum processing units, and the Maximum processing units. How many processing units you require is based entirely upon your network traffic.

Note that processing units are measured in whole processors. Setting the Maximum processing units to 1 would mean that at most, the LPAR you are creating will use one entire processor. We have set the Desired processing units to 0.5, the Minimum processing units to 0.1, and the Maximum processing units to 1.0, as shown in Figure 2-7 on page 27.

Important: The security tools partition will likely be doing more processing than the firewall partition, and therefore may require more processing units to be assigned to it.

Once you have selected your desired processing unit level, click the **Next** button to continue.




Figure 2-7 LPAR - Processing Settings

- h. Assign physical hardware to the LPAR.

Note: As shown in Figure 2-1 on page 21, the LINUXFW partition requires two physical Ethernet adapters to be assigned to it. All other partitions are connected to each other via virtual Ethernet connections, and do not require any physical cards to be linked to them. Do not add any physical cards to the LINUXST partition during setup.

To link two physical Ethernet cards to the LINUXFW partition, follow these steps:

- i. Expand the Managed system I/O pool that you wish to use.
- ii. Expand the bus that two available Ethernet adapters are located on.
- iii. Highlight two available Ethernet adapters by holding the Ctrl key and clicking each desired adapter once.
- iv. With the Ethernet adapters still highlighted, click the **Add as required** button.

The devices will appear under the Profile I/O devices heading, as shown in Figure 2-8 on page 28. Click the **Next** button to continue.




Figure 2-8 LPAR - Specifying physical Ethernet adapters

- i. Edit the partition's I/O pool participation, as shown in Figure 2-9 on page 29. If you have I/O pools created on your eServer i5, you may wish to have this LPAR draw resources from them. If you do not have I/O pools defined, no action is required.

When you have finished specifying any required I/O pool options, click the **Next** button.




Figure 2-9 LPAR - Specifying I/O pool participation

- j. Specify required virtual I/O adapters. Select **Yes, I want to specify virtual I/O adapters**, as shown in Figure 2-10 on page 30, and click the **Next** button to continue.

Attention: Your eServer i5 must have PTF MF33433 installed. If not, virtual devices cannot be used.




Figure 2-10 LPAR - Virtual I/O Adapters

- k. Create virtual I/O adapters. The LPARs that you are creating will not have physical hard disks of their own. They will also lack physical serial connections for a console to use for setup and administration. As such, it is necessary to utilize the virtual I/O features of the eServer i5 to create virtual SCSI and serial adapters for the Linux partitions to use. In addition, communication between the partitions will be accomplished via virtual Ethernet connections.

Two server serial virtual I/O devices should exist by default, as shown in Figure 2-11 on page 31. Your new LPAR requires two more virtual I/O devices: One virtual Ethernet adapter and one virtual SCSI adapter.




Figure 2-11 LPAR - Create Virtual I/O Adapters

To create the required virtual Ethernet adapters, follow these steps:

- i. Select the **E**thernet radio button under the Create adapters heading.
- ii. Click the **C**reate button.
- iii. The Virtual Ethernet Adapter Properties window will appear, as shown in Figure 2-12 on page 32. Leave the default values as they are. Click the **O**K button to create the adapter. The adapter appears in list of virtual I/O devices.
- iv. Click the **R**equired check box.
- v. If you are creating the LINUXFW partition, repeat this step. The LINUXFW partition requires two virtual networks so that it can communicate with both the i5/OS partition and the LINUXST partition.




Figure 2-12 LPAR - Virtual Ethernet Adapter Properties

To create the required virtual SCSI adapters, follow these steps:

- i. Select the **SCSI** radio button.
- ii. Click the **Create** button.
- iii. The Virtual SCSI Adapter window appears.
- iv. Leave the Slot number at the default setting. Remember this number, as you will use it to create the virtual SCSI server adapter on the i5/OS partition.
- v. Ensure that the Adapter Type is set to Client.
- vi. The Remote partition should be the i5OS partition.
- vii. Enter 0 for the Remote slot number. This will need to be changed later. This cannot be done until the virtual SCSI server is created on the i5/OS partition.
- viii. Click the **OK** button to create the virtual SCSI client adapter.
- ix. Check the **Required** check box beside the SCSI adapter you just created.

After adding all virtual I/O devices, your window should look similar to Figure 2-13 on page 33. Click the **Next** button to continue.




Figure 2-13 Created virtual I/O adapters

5. Set up power controlling partitions. We recommend that the i5/OS partition is the only partition that you grant authority to control power for the LINUXFW and LINUXST partitions.

Ensure that **i5OS** is selected in the Power controlling partition from the drop-down and press the **Add** button. The i5/OS partition is added to the list of power controlling partitions, as shown in Figure 2-14 on page 34. Click the **Next** button to continue.




Figure 2-14 LPAR - Power Controlling Partitions

- I. Set optional LPAR settings. We recommend that you do *not* select the Automatically start with managed system check box, as the LPAR uses virtual I/O.

For now, you *must* select the **System Management Services (SMS)** radio button under the Boot modes heading, as shown in Figure 2-15 on page 35. This will allow registration of the LPAR to the Open Firmware on the system. Click the **Next** button to continue.




Figure 2-15 LPAR - Optional Settings

- m. Read the profile summary that you are presented with, as shown in Figure 2-16 on page 36. Ensure that the displayed information is accurate before clicking the **Finish** button to create the LPAR. Be patient while the LPAR is created.




Figure 2-16 LPAR - Summary

6. Activate the partition. In the Server and Partition: Server Management pane, right-click the newly created LPAR. Select **Activate**, as shown in Figure 2-17 on page 37.




Figure 2-17 LPAR - Activation

7. The Activate Logical Partition window appears.
 - a. Ensure that the **Open a terminal window or console session** check box is selected, as shown in Figure 2-18.
 - b. Click the **OK** button to activate the partition.




Figure 2-18 LPAR - Activate Logical Partition

8. The SMS console appears.
 - a. If you receive the message shown in Figure 2-19, press the 0 key.




Figure 2-19 LPAR - SMS console

- b. When you see the screen shown in Figure 2-20 on page 39, the LPAR has been successfully booted and registered. You may close the console window.

```
Version SF220_051
SMS 1.5 (c) Copyright IBM Corp. 2000,2003 All rights reserved.
-----
Main Menu
 1. Select Language
 2. Setup Remote IPL (Initial Program Load)
 3. Change SCSI Settings
 4. Select Console
 5. Select Boot Options

-----

Navigation Keys:

                                X = eXit System Management Services
-----

Type the number of the menu item and press Enter or select Navigation Key:
```

Figure 2-20 LPAR - Successful SMS boot

9. Back in the Server and Partition: Server Management pane, right-click the partition you just created, and select **Shut Down Partition**, as shown in Figure 2-21 on page 40, to shut down the running LPAR.




Figure 2-21 LPAR - Shut down

10. The Shut Down Partitions window appears.

- a. Select the **Immediate** radio button as the shut down option.
- b. Click the **OK** button. The server will shut down.

This completes the creation of an LPAR.

2.2.2 Set up the i5/OS partition virtual I/O

The Linux partitions that you have just created do not have physical hard disks of their own. They also lack physical serial connections for a viewing console. As such, it is necessary to use the virtual I/O features of the eServer i5 to create virtual SCSI adapters and virtual consoles for the Linux partitions to use.

Setting up virtual SCSI server adapters

Virtual SCSI adapters allow LPARs to read and write to any disk pool that is managed by the i5/OS. For the adapters to function, there must be a server adapter created on the i5/OS, and a client adapter created on the LPAR. Now we will create virtual SCSI server adapters. Creating client adapters is covered in 2.2.1, “Creating a logical partition using the HMC” on page 20. To create virtual SCSI adapters, complete the following steps.

Important: The following steps create one virtual SCSI server adapter. It is imperative that these steps be repeated so that two virtual SCSI adapters will be created. On the first iteration of the steps, use the values in column 2 on the planning worksheet to set up the firewall partition. On the second iteration of the steps, use the values from column 3 to set up the security tools partition.

1. In the Server and Partition: Server Management pane, expand i5OS.
2. Right-click the profile that you use to run your i5/OS partition, and select **Properties**, as shown in Figure 2-22.




Figure 2-22 i5/OS partition properties

3. The Logical Partition Profile Properties window appears.
 - a. Click the **Virtual I/O** tab.
 - b. Select the **SCSI** radio button.
 - c. Click the **Create** button.
 - d. The Virtual SCSI Adapter window appears.
 - i. Leave the Slot number at the default setting. Remember this number, as this is the value that you will need to enter for the Remote partition client SCSI adapter in the virtual I/O settings.
 - ii. Ensure that the Adapter Type is set to Server.
 - iii. The Remote partition should be the appropriate value from row B on the planning worksheet.
 - iv. The Remote partition virtual slot number should be set to the number that you noted in step iv on page 32. The settings should look similar to those shown in Figure 2-23 on page 42.




Figure 2-23 i5/OS virtual SCSI server adapter settings

- v. Click the **OK** button to save the settings.
- e. Select the **Required** check box beside the SCSI adapter you just created.
- f. Click the **OK** button on the Logical Partition Profile Properties window.
4. Back in the main HMC window, in the Server and Partition: Server Management pane, expand the appropriate LPAR name from row B of the planning worksheet.
5. Right-click the **Default** profile and select **Properties**.
6. The Logical Partition Profile properties window appears.
 - a. Select the **Virtual I/O** tab.
 - b. Click the **Client SCSI** adapter that you created on page 32.
 - c. Click the **Properties** button.
 - d. The Virtual SCSI Adapter Properties window appears.
 - i. Enter the number that you noted in step iv on page 32 for the Remote partition virtual slot number.
 - ii. Click the **OK** button to save the virtual SCSI server adapter settings.
 - e. Click the **OK** button on the Logical Partition Profile Properties window.

The virtual SCSI server adapter has been created.

Setting up virtual serial client adapters

Virtual serial adapters allow you to view the output of a Linux console running on an LPAR through an Ethernet connection, rather than a physical serial connection to the Linux server. To create virtual serial adapters, complete the following steps.

Important: The following steps will create one virtual serial server adapter. It is imperative that these steps be repeated so that two virtual serial adapters will be created. On the first iteration of the steps, use the values in column 2 on the planning worksheet to create the firewall partition. On the second iteration of the steps, use the values from column 3 to create the security tools partition.

1. In the Server and Partition: Server Management pane, expand i5OS.
2. Right-click the profile that you use to run your i5/OS partition, and select **Properties**.
3. The Logical Partition Profile Properties window appears.
 - a. Click the **Virtual I/O** tab.
 - b. Select the **Serial** radio button.
 - c. Click the **Create** button.
 - d. The Virtual Serial Adapter window appears.
 - i. Leave the Slot number at the default setting.
 - ii. Ensure that the Adapter Type is set to Client.
 - iii. The Remote partition should be the appropriate value from row B of the planning worksheet.
 - iv. The Remote partition virtual slot number should be set to 0. Your settings should look like those shown in Figure 2-24.




Figure 2-24 i5/OS virtual serial client adapter settings

- v. Click the **OK** button to save the settings.
- e. Click the **OK** button on the Logical Partition Profile Properties window.

The virtual serial client adapter has been created.

Attention: The i5/OS partition *must* be restarted for these virtual I/O changes to be activated. The Linux operating system installation cannot proceed until the i5/OS partition is restarted.

2.2.3 Working with network servers

Before the Linux operating system can be installed on the newly created LPAR, an NWSSTG and NWSD must be created, and a link must be formed between them.

Important: The following instructions show how to create and link the NWSSTG and NWSD for one partition. These steps *must* be repeated to create the second NWSSTG and NWSD. On the first iteration of the steps, use the values in column 2 on the planning worksheet to set up the firewall partition. On the second iteration of the steps, use the values from column 3 to set up the security tools partition.

Creating the NWSSTG

To create the required Network Server Storage Space, follow these steps:

1. Open iSeries Navigator, and expand your server name and then select **Network** → **Windows Administration**.
2. Right-click **Disk Drives** and select **New Disk**, as shown in Figure 2-25.




Figure 2-25 Create a new NWSSTG

3. The New Disk window appears. Specify the Disk drive name as the value from row F of the planning worksheet.
4. Specify the Description as the value from row G of the planning worksheet.

- Specify the Capacity as the value from row H of the planning worksheet.

Note: The LINUXST partition may require more space than 8 GB, depending on the volume of mail you receive, and the level of logging you decide to enable on your system. If you receive high volumes of mail, or enforce maximum logging policies, you may decide to increase the capacity.

- Your settings should look like Figure 2-26. Click the **OK** button to begin creation of the NWSSTG.



Figure 2-26 New disk settings

- An animation will play while the disk is created, as shown in Figure 2-27. This can take several minutes. Once the disk has been successfully created, the window will close automatically and the disk will be shown in the iSeries Navigator window.



Figure 2-27 Creating the NWSSTG

The NWSSTG has been created. Now you must create the NWSD.

Creating the NWSD

To create the required Network Server Storage Space, follow these steps:

- At a 5250 command line to your eServer i5, enter the CRTNWSD command, and press F4.
- Press F9. You will be presented with the screen shown in Figure 2-28 on page 46.

```

Create Network Server Desc (CRTNWSD)

Type choices, press Enter.

Network server description . . . . . Name
Resource name . . . . . Name, *NONE, *AUTO
Network server type . . . . . *WINDOWSNT *WINDOWSNT, *GUEST
Online at IPL . . . . . *YES *YES, *NO
Vary on wait . . . . . *NOWAIT *NOWAIT, 1-15 minutes
Shutdown timeout . . . . . 15 2-45 minutes
Partition . . . . . *NONE

Partition number . . . . . *NONE Number, *NONE
Domain role . . . . . *DMNCTL *DMNCTL, *BKUCTL, *SERVER
Propagate domain users . . . . . *YES *YES, *NO
Language version . . . . . *PRIMARY *PRIMARY, 2963, 2966, 2980...
Code page . . . . . *LNGVER *LNGVER, 437, 850, 852, 857...
Server message queue . . . . . *JOBLOG Name, *JOBLOG, *NONE
Library . . . . . Name, *LIBL, *CURLIB

More...

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

```

Figure 2-28 CRTNWSD screen

3. Enter the appropriate value from row D on the planning worksheet as the Network server description.
4. For Resource name, enter *AUTO.
5. Enter *GUEST as the Network server type.
6. Change Online at IPL to *NO.
7. Change Partition to the appropriate value from row B of the planning worksheet.
8. Change the value of Code page to 437. The first page of settings should look similar to Figure 2-29 on page 47.

```

Create Network Server Desc (CRTNWS D)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . *SYS
+ for more values

Ports:
Port number . . . . . *NONE      1, 2, *INTERNAL, *VRTETHPTP...
Line description . . . . .      Name
+ for more values
Restricted device resources . . *NONE      Name, *NONE, *ALL...
+ for more values
Synchronize date and time . . . *TYPE      *TYPE, *YES, *NO
IPL source . . . . . *NWSSTG      *NWSSTG, *PANEL, *STMF, A...
IPL stream file . . . . . *NONE

More...

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

```

Figure 2-29 Create network server description settings

9. Page down three times. Change the value of IPL source to *STMF.
10. Change the value of IPL stream file to '/QOPT/SU90.001/INSTALL'.
11. Your settings page should look similar to Figure 2-30.

```

Create Network Server Desc (CRTNWS D)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . *SYS
+ for more values

Ports:
Port number . . . . . *NONE      1, 2, *INTERNAL, *VRTETHPTP...
Line description . . . . .      Name
+ for more values
Restricted device resources . . *NONE      Name, *NONE, *ALL...
+ for more values
Synchronize date and time . . . *TYPE      *TYPE, *YES, *NO
IPL source . . . . . > *STMF      *NWSSTG, *PANEL, *STMF, A...
IPL stream file . . . . . > '/QOPT/SU90.001/INSTALL'

More...

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

```

Figure 2-30 More create network server description settings

12. Press the Enter key to create the NWS. You will see the message Network server description LINFWS created at the bottom of the screen.

The NWS has been created. Now you must simply link the NWSSTG to the NWS before proceeding with the installation of your Linux operating system.

Creating the required network sever link

To create the required network server link:

1. Open iSeries Navigator and expand your server name, then click **Network** → **Windows Administration**.
2. Click **Disk Drives**.
3. A list of NWSSTGs that are currently on your i5 appears on the right, as shown in Figure 2-31.




Figure 2-31 Existing NWSSTGs

4. Right-click the appropriate NWSSTG name from row F of the planning worksheet, and select **Add Link**, as shown in Figure 2-32 on page 49.




Figure 2-32 Linking the NWSSTG to the NWSD

5. The Add Link to Server window appears.
 - a. Select the appropriate NWSD from row D of the planning worksheet from the list of NWSDs for the Server to link to.
 - b. Leave all other settings at their default. Your settings should look similar to Figure 2-33.




Figure 2-33 Link settings

- c. Click the **OK** button. The window will disappear, and the link will be created.

You have now created an NWSSTG that uses the appropriate virtual SCSI adapter, and an NWSM that is linked to the NWSSTG. You may now proceed to installing your Linux operating system on your LPAR.

2.3 Installing Linux

Now it is time to install your Linux operating system on your newly created LPAR. This section covers the installation of SUSE Linux Enterprise Server 9 (SLES9) and RedHat Enterprise Linux 4 (RHEL4).

2.3.1 Installing SLES9

Note: The following steps give instructions for installing the SLES9 operating system on one of the LPARs you have just created. You must repeat these steps to install SLES9 on both of the new LPARs. On the first iteration of the steps, use the values in column 2 on the planning worksheet to set up the firewall partition. On the second iteration of the steps, use the values from column 3 to set up the security tools partition.

To install Linux on your LPAR, follow these steps:

1. Insert the first disk from your Linux distribution into the CD-ROM drive on your eServer i5.
2. Open the virtual console connection:
 - a. Open up the PuTTY telnet client.
 - b. Enter your i5/OS server name as the Host Name (or IP address).
 - c. Set the protocol to Telnet.
 - d. Enter 2301 as the Port. Your settings should look similar to Figure 2-34 on page 51. Press the **Open** button to open the connection.




Figure 2-34 PuTTY settings

- e. You will be presented with a screen similar to the one shown in Figure 2-35. Enter the number that corresponds with the appropriate NWSD name from row D on the planning worksheet, and press the Enter key.

```
OS/400 Guest Partition Consoles
2: Linux_ST(V1-C2/V2-C0)
3: Linux_FW(V1-C4/V3-C0)
```

Enter the console partition number:

Figure 2-35 OS/400 Guest Partition Consoles

- f. You will be asked to enter your service tools user ID. Enter it now, and press the Enter key.
- g. Enter your service tools user ID and password, and press the Enter key. You will see the screen shown in Figure 2-36. Leave this PuTTY terminal open.

```
Linux_FW: Enter OS/400 service tools userid:
linuxusr
Linux_FW: Enter OS/400 service tools password:
Linux_FW: Console connecting...
Linux_FW: Console connected.
```

Figure 2-36 Console connected

3. Set the NWSD boot parameters:
 - a. Open a 5250 command line to your eServer i5.
 - b. Enter the WRKNWSD command, and press the Enter key.
 - c. Place a 2 beside LINFWSD and press the Enter key.
 - d. Page down twice until you see the screen in Figure 2-37.

```
Change Network Server Desc (CHGNWSD)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . *SYS
+ for more values
Restricted device resources . . *NONE      Name, *SAME, *NONE, *ALL...
+ for more values
Synchronize date and time . . . *NO      *SAME, *TYPE, *YES, *NO
IPL source . . . . . *STMF      *SAME, *NWSSTG, *PANEL...
IPL stream file . . . . . *NONE

IPL parameters . . . . . *NONE

More...

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

Figure 2-37 Change NWSD

- e. Change the value of IPL stream file to ‘/QOPT/SU90.001/INSTALL’.
- f. Press the Enter key to save the changes. You will see the message Description for network server changed printed at the bottom of the screen.
- g. Press the F3 key.
4. Vary on the NWSD:
 - a. Type WRKCFGSTS (*NWS) at the 5250 command line, and press the Enter key.

- b. You will be shown a list of NWSDs currently on your system, along with their status, as shown in Figure 2-38.

```

Work with Configuration Status                                RCHAS10
                                                            04/25/05 10:12:06
Position to . . . . . Starting characters
Type options, press Enter.
 1=Vary on  2=Vary off  5=Work with job  8=Work with description
 9=Display mode status 13=Work with APPN status...

Opt  Description      Status      -----Job-----
    LINFWSD           VARIED OFF
    LINSTSD           VARIED OFF

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

Bottom

```

Figure 2-38 Work with Configuration Status screen

- c. Type a 1 beside the appropriate NWSD name from row D on the planning worksheet, and press Enter.
 - d. The status of the NWSD will change to VARY ON PENDING, and the message Vary on completed for network server will be printed at the bottom of the screen.
 - e. If you refresh the screen by pressing the F5 key, the status should read ACTIVE. Do not end this 5250 session; leave it open.
5. Back in the PuTTY terminal that you opened earlier, you should now begin to see output from the LPAR boot sequence. The first boot messages you see should look like the screen in Figure 2-39 on page 54.


```
YaST                                         Press F1 for Help
+-----+
|SUSE LINUX Enterprise Server (SLES) 9      -
|Novell Software License Agreement         -
|
|PLEASE READ THIS AGREEMENT CAREFULLY. BY INSTALLING OR
|OTHERWISE USING THE SOFTWARE (INCLUDING ITS COMPONENTS), YOU
|AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE
|WITH THESE TERMS, DO NOT DOWNLOAD, INSTALL OR USE THE
|SOFTWARE AND, IF APPLICABLE, RETURN THE ENTIRE UNUSED
|PACKAGE TO THE RESELLER WITH YOUR RECEIPT FOR A REFUND. THE
|SOFTWARE MAY NOT BE SOLD, TRANSFERRED, OR FURTHER
|DISTRIBUTED WITHOUT PRIOR WRITTEN AUTHORIZATION FROM NOVELL.
|
|RIGHTS AND LICENSES
|
|This Novell Software License Agreement ("Agreement") is a
|legal agreement between You (an entity or a person) and
|Novell, Inc. ("Novell") with respect to the software
|product identified in the title of this Agreement, media (if
+-----+
|                                     [ I Agree ]
|                                     [I Do Not Agree]
```

Figure 2-40 YaST - Novell Software License Agreement

7. Use the Tab key to select the **I Agree** button, and press the Enter key.
8. The language selection screen in Figure 2-41 on page 56 is shown. Use the arrow keys to select your language from the list. Tab to the Accept button and press the Enter key.




Figure 2-41 YaST - Language selection

9. You will be presented with an error that informs you that the parted tool cannot read your hard disk, as shown in Figure 2-42. This is normal, and will not affect your installation. Click the **OK** button.




Figure 2-42 YaST - Parted error

10. You will be presented with a suggested setup for your new Linux system, as shown in Figure 2-43 on page 57. This setup contains packages that we do not require, so they must first be removed before the installation proceeds. Change the installation type to the minimum installation:
 - a. Tab to the Change button and press the Enter key.

```

YaST                                                    Press F1 for Help

+-----+ Installation Settings
| Use Accept to - Click any headline to make changes or use the "Change
| perform a new +-----+
| installation with |System -
| the values      | * System: CHRP
| displayed.     | * Processor: 2x POWER5 (gr)
| Change the values | * Main Memory: 2 GB
| by clicking on the |Mode
| respective headline | * New installation
| or by using the |Keyboard layout
| Change... menu.   | * English (US)
| To update an |Mouse
| existing SUSE Linux | * NONE
| system instead of |Partitioning
| doing a new install, | * Create boot partition 17.0 MB on /dev/sda1
| click the Mode - | * Create swap partition 998.0 MB on /dev/sda2
| headline or select | * Create root partition 7.0 GB (/dev/sda3 with
| Mode in the      | reiser)
| Change... menu. +-----+
| Your hard disk   |
+-----+
                                     [Change...-]
                                     [Abort]
                                     [Accept]

```

Figure 2-43 YaST - Suggested setup

- b. Select **Software** from the list that pops up, and press the Enter key.
- c. Tab to the Minimum System option, and use Spacebar to toggle the radio button associated with it, as shown in Figure 2-44 on page 58.




Figure 2-44 YaST - Software selection

- d. Tab to the Accept button, and press the Enter key.
11. The system setup values have now been changed. Tab to the Accept button and press the Enter key.
12. You will receive the warning shown in Figure 2-45. Tab to Yes, install and press the Enter key.




Figure 2-45 YaST - Ready to continue warning

13. YaST will prepare your hard disks, as shown in Figure 2-46 on page 59.

```
YaST                                     Press F1 for Help
+-----+ Preparing Your Hard Disk
| Please wait while-
| your hard disk is
| prepared for
| installation...
| Depending on the
| size of your hard
| disk and your
| processor speed,
| this action might
| take some time. 5
| minutes are not
| unusual for disks
| larger than 4 GB.
| Often, the progress
| meter doesn't show a-
| linear progress;
| even if it looks
| slow near the end
| ("95 %"), please be
+-----+
|                                     Formatting /dev/sda3 as /
|                                     64%
|
|                                     [Abort]           [Accept]
```

Figure 2-46 YaST - Preparing disk drives

14. The installation will begin. When you are asked to switch disks, insert the requested CD, and press the Enter key. You may have to press the Enter key twice before YaST accepts the new CD.
15. Installation will proceed, as shown in Figure 2-47 on page 60. You will be updated as the installation progresses. When installation has completed, the server will reboot.

```

YaST                                     Press F1 for Help
[Back]                                [Abort Installation]                                [Next]
+Current Package-----+ | Time remaining |
| coreutils-5.2.1 | | (estimated) |
|                | | SUSE SLES |
|                | | CD 1: [x] |
|                | | 0% |
+-----+ |
|Current Package | | SUSE CORE |
| * coreutils-5.2.1 | | CD 1: 07:47 |
|Description | | 35% |
| * GNU Core Utilities | |
+-----+ |
|                | | CD 2: --- |
+-----+ |
|                | | CD 3: --- |
|                | | CD 4: --- |
+-----+ |
|                | | CD 5: --- |
+-----+ |
|Installation Log (Extract)-----+ | Total: 15:55 |
|e2fsprogs-64bit-9 --- Utilities for the second ex |
|~tended file system | |
|coreutils-5.2.1 --- GNU Core Utilities | |
+-----+ | [Slide Show...] |

```

Figure 2-47 YaST - Installation progresses

16. Vary off the LPAR:

- a. Back in the 5250 screen that you opened earlier, place a 2 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- b. The status of the NWSD will change to VARY OFF PENDING. It will take several minutes before the NWSD is varied off. Occasionally press the F5 key to refresh the screen. When the status of the NWSD reads VARIED OFF, proceed to the next step.
- c. Press the F3 key exit the Work with Configuration Status screen.

17. Change the LPAR boot options to boot from the NWSSTG:

- a. At the 5250 command line, enter WRKNWSD and press the Enter key.
- b. Place a 2 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- c. Page down twice so that you have the IPL options on your screen.
- d. Change IPL source to *NWSSTG.
- e. Change IPL stream file to *SAME.
- f. Change IPL parameters to 'root=/dev/sda3'. Your options should look like Figure 2-48 on page 61.

```

Change Network Server Desc (CHGNWSD)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . *SYS
+ for more values
Restricted device resources . . *NONE      Name, *SAME, *NONE, *ALL...
+ for more values
Synchronize date and time . . . *NO      *SAME, *TYPE, *YES, *NO
IPL source . . . . . *NWSSTG      *SAME, *NWSSTG, *PANEL...
IPL stream file . . . . . *SAME

IPL parameters . . . . . 'root=/dev/sda3'

More...

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

```

Figure 2-48 Changed NWSD IPL settings

- g. Press the Enter key to save the changed NWSD settings. The message Description for network server changed will be printed at the bottom of your screen.
 - h. Press the F3 key to exit the Change Network Server Description screen.
18. Change the LPAR profile boot options:
- a. In the main HMC window, expand the appropriate LPAR name from row B of the planning worksheet.
 - b. Right-click the **Default** profile, and select **Properties**.
 - c. The Logical Partition Profile Properties window appears.
 - i. Click the **Settings** tab.
 - ii. Under the Boot Modes heading, select **Normal**, as shown in Figure 2-49 on page 62.




Figure 2-49 Changing boot modes

iii. Click the **OK** button to save the settings.

19. Vary on the NWSD:

- a. At a 5250 command line to your eServer i5, type `WRKCFGSTS (*NWS)` and press the Enter key.
- b. You will be shown a list of NWSDs currently on your system, along with their status. Type a 1 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- c. The status of the NWSD will change to `VARY ON PENDING`, and the message `Vary on completed for network server` will be printed at the bottom of the screen.
- d. If you refresh the screen by pressing the F5 key, the status should read `ACTIVE`. Do not end this session; leave it open.

20. Back in your PuTTY terminal window, you can see your Linux operating system booting up. The boot process is not interactive, and requires no input from the user.

21. After installation, SLES9 will return to YaST on the initial boot of the system. This is so that you can set the root user password, and set up your Ethernet cards. Type the password you would like your root user to have.

22. Tab to the next field, and re-enter the same password, as shown in Figure 2-50 on page 63.

```

YaST @ linux                                     Press F1 for Help

+-----+ Password for "root", the system administrator
| Unlike normal | -
| users of the system, |
| who, for instance, |
| write texts, create |
| graphics, or browse |
| the Internet, the | -
| user "root" exists |
| on every system and |
| is called into |
| action whenever |
| administrative tasks |
| need to be |
| performed. Log in as |
| root when you need |
| to be the system |
| administrator and |
| only then. |
| Because the root |
| user is equipped |
+-----+
| [Back] | [Abort] | [Next]

```

Figure 2-50 YaST - Setting the root password

23. Tab to the Next button, and press the Enter key to set the root password.
24. You are now shown the Network Configuration screen, as shown in Figure 2-51 on page 64. This screen allows you to configure all physical and virtual network interfaces that enable your server to communicate with the Internet and your LAN, as well as with other partitions on the eServer i5. Configure your network interfaces:
 - a. Tab to the Change button, and press the Enter key.

```

YaST @ linux                                     Press F1 for Help

+-----+ Network Configuration
| Put the network settings into effect by pressing Next. Change the values by clicking on the respective headline or by using the Change... menu. |
+-----+
| ( ) Skip Configuration |
| (x) Use Following Configuration |
+-----+
| Network Interfaces |
| * IBM 79c970 [PCnet32 LANCE] |
|   Configured with DHCP |
| * IBM 79c970 [PCnet32 LANCE] |
|   Not configured yet. |
| * IBM Virtual Ethernet card 1 |
|   Not configured yet. |
| * IBM Virtual Ethernet card 2 |
|   Not configured yet. |
| Proxy |
| * Proxy is disabled. |
| VNC Remote Administration |
+-----+
| [Change...-] |
+-----+
| [Back] | [Abort] | [Next]

```

Figure 2-51 YaST - Network configuration

- b. Select **Network Interfaces** from the pop-up list, and press the Enter key. You will be presented with the Network cards configuration screen, as shown in Figure 2-52.

```

YaST @ linux                                     Press F1 for Help

+-----+ Network cards configuration
| Network card setup- |
| Configure your network card here. Adding a network card: Choose a network card from the list of detected network cards. If your network card was not autodetected, select Other (not detected) then press Configure. |
| Editing or Deleting: If you press Change, an additional dialog |
+-----+
| Network cards to configure-----+
| +Available are:-----+
| | IBM 79c970 [PCnet32 LANCE] |
| | IBM Virtual Ethernet card 1 |
| | IBM Virtual Ethernet card 2 |
| | Other (not detected) |
+-----+
| [Configure...] |
+-----+
| Already configured devices: |
| * IBM 79c970 [PCnet32 LANCE] |
|   Configured with DHCP |
+-----+
| [Change...] |
+-----+
| [ Back ] | [Abort] | [Finish]

```

Figure 2-52 YaST - Network cards configuration

- c. All of the Ethernet connections must have statically assigned IP addresses. Therefore the first task is to remove the adapter configuration that was automatically set to use DHCP.
 - i. Tab to the Change button and press the Enter key. You will be presented with the Network card configuration overview, shown in Figure 2-53.

```

YaST @ linux                                     Press F1 for Help

+-----+ Network cards configuration overview
| Network card overview |
| Obtain an overview of |
| installed network cards. |
| Additionally, edit their |
| configuration. |
| Adding a network card: |
| Press Add to configure a new |
| network card manually. |
| Editing or deleting: |
| Choose a network card to change or |
| remove. Then press Edit or Delete as |
+-----+

+-----+
| Name | Device | IP |
+-----+
| IBM 79c970 [PCne... | eth-id-00:09:6b:65:37:62 | DH |
+-----+

[Add] [Edit] [Delete]

[ Back ] [Abort] [Finish]

```

Figure 2-53 YaST - Network cards configuration overview

- ii. Tab to the Delete button and press the Enter key.
- iii. Tab to the Finish button and press the Enter key.
- d. You now have a completely blank network configuration. Configure your static addresses:
 - i. Tab to the Change button and press the Enter key.
 - ii. Select **Network interfaces** and press the Enter key.
 - iii. Tab to the Available Are selection box, and use the arrow keys to select the first IBM Virtual Ethernet card.
 - iv. Tab to the Configure button and press the Enter key.
 - v. Enter the value for the Configuration Name from row AA on the planning worksheet.
 - vi. Tab to Static Address setup and press the Spacebar to select it.
 - vii. Tab to the IP Address text box and enter in the value from row AB on the planning worksheet.
 - viii. Tab to the Subnet mask text box and enter in the value from row AC on the planning worksheet. Your settings should look similar to those in Figure 2-54 on page 66.

```

YaST @ linux                                     Press F1 for Help

+-----+ Network address setup
| Configure your IP -
| address.
| You can select
| dynamic address
| assignment, if you
| have a DHCP server
| running on your
| local network.
| Also select this
| if you do not have a
| static IP address
| assigned by the
| system administrator
| or your cable or DSL
| provider.
| Network addresses
| will then be
| obtained
| automatically from
+-----+
| [Back]
+-----+

Configuration Name
eth2|

+Choose the setup method-----+
| ( ) Automatic address setup (via DHCP)
| (x) Static address setup
| IP Address      Subnet mask
| 10.1.1.10|      255.255.255.128|
+-----+

+Detailed settings-----+
| [Host name and name server]
| [ Routing ]
| [ Advanced... -]
+-----+

| [Back]
| [Abort]
| [Next]
+-----+

```

Figure 2-54 YaST - Network address setup

- e. Configure host name and DNS settings:
 - i. Tab to the Host name and name server button and press the Enter key.
 - ii. Enter the appropriate value for the Host Name from row N of the planning worksheet.
 - iii. Enter the appropriate value for the Domain Name from row O of the planning worksheet.
 - iv. Enter in your name server information from rows L and M on the planning worksheet.
 - v. Enter in the appropriate value for Domain Search 1 from row O of the planning worksheet. Your settings should look like Figure 2-55 on page 67.

```

YaST @ linux                                     Press F1 for Help

+-----+ Host name and name server configuration
| Insert the host -                               |
| name and domain name|                         |
| for your computer. | +Host name and domain name-----+
| Name server list and| |Host Name           Domain Name   | |
| domain search list | |LinuxFW|            rchland.ibm.com|
| are optional.      | | [ ] Change host name via DHCP  |
| A name server is a- +-----+
| computer that      |
| translates host    | +Name servers and domain search list-----+
| names into IP      | |Name Server 1       Domain Search 1  | |
| addresses. This    | |9.10.244.200|       rchland.ibm.com|
| value must be      | |Name Server 2       Domain Search 2  |
| entered as an IP   | |9.10.244.100|       |
| address (e.g.,     | |Name Server 3       Domain Search 3  |
| 10.10.0.1), not as a| |
| host name.         | | [ ] Update name servers and search list via DHCP|
| Search domain is   +-----+
| the domain name    |
| where host name    |
+-----+ [Back]                [Abort]                [ OK ]

```

Figure 2-55 YaST - Host name and name server configuration

- vi. Tab to the OK button and press the Enter key.
- f. Set up your default gateway:
 - i. Tab to the Routing button and press the Enter key.
 - ii. Enter in the value for the Default Gateway. If you are setting up the LINUXFW partition, this value should be the gateway for your LAN (cell K2 on the planning worksheet). If you are setting up the LINUXST partition, this should be the virtual LAN IP address of the LINUXFW partition (cell AB2 on the planning worksheet).
 - iii. Tab to the OK button and press the Enter key.
- g. Tab to the Next button and press the Enter key.
- h. If you are installing the LINUXFW partition, you will need to configure each remaining Ethernet card. Tab to the Available are selection box, select the first detected network card in the list, tab to Configure, and press the Enter key. Configure eth0 using cells CA2 through CC2, and configure eth1 using cells BA2 through BC2 from the planning worksheet.
- i. When all required Ethernet adapters have been configured, tab to the Finish button and press the Enter key.
- j. Tab to the Next button and press the Enter key to write out the network configuration to the hard disk. You will see the screen shown in Figure 2-56 on page 68.

28. Add a local user to the system. Fill in the required values. Your settings should look similar to Figure 2-60 when you are finished. Tab to the Next button and press the Enter key.




Figure 2-60 YaST - Add a New Local User

29. All of the settings you have just entered will now be saved. You will be shown release notes for SLES9 on IBM POWER™. Tab to the Next button and press the Enter key.

30. You will be presented with a success message, shown in Figure 2-61 on page 71. Tab to the Finish button and press the Enter key. The system will boot into your newly installed Linux operating system.


```

YaST @ LinuxFW                                     Press F1 for Help
+-----+
|                                     YaST Control Center                                     |
+-----+
+-----+-----+-----+-----+
| Software      | | Online Update |
| Hardware     | | Install and Remove Software |
| System       | | Change Source of Installation |
| Network Devices | | Installation into Directory |
| Network Services | | Patch CD Update |
| Security and Users | | System Update |
| Misc         | | YOU Server Configuration |
+-----+-----+-----+-----+
[Help]                                               [Quit]

```

Figure 2-62 YaST - Main screen

- d. Tab to the next window and select Install and Remove Software. Press the Enter key.
- e. Press Alt+F to load the Filter menu.
- f. Use the arrow keys to select Selections, as shown in Figure 2-63, and press the Enter key.

```

YaST @ LinuxFW                                     Press F1 for Help
[Filter-]      [Actions-]      [Information-]      [Etc.-]
+-----+-----+-----+-----+
| RPM Groups   | | Avail. Vers. | | Inst. Vers. | | Summary | |
| Selections  | | 1.31         | |              | | X Window System background |
| Search      | | 4.16         | |              | | A screen saver and locker |
| Installed Packages | | me|4.16     | |              | | Gnome bindings for xscreen |
| Installation Summary |
| Update List |
+-----+-----+-----+-----+
Filter: Amusements                               Required Disk Space: 0 B
+-----+-----+-----+-----+
|xbanner - X Window System background writings and images
|Version: 1.31-858.1 Size: 501.1 kB Media No.: 1
|License: GPL, Other License(s), see package
|Package Group: Amusements/Toys/Background
|Provides:
|Authors: Amit Margalit <amitm@netvision.net.il>
+-----+-----+-----+-----+
[Help-]      [Search]      [Disk Usage]      [Cancel]      [Accept]

```

Figure 2-63 YaST - Filter menu

- g. Locate C/C++ Compiler Tools in the list. Select that entry, and press the Enter key. A plus will appear beside the entry in the list, as shown in Figure 2-64.




Figure 2-64 YaST - Selecting C/C++ Compiler Tools

- h. Tab to the OK button and press the Enter key.
- i. Press Alt+F to load the Filter menu.
- j. Use the arrow keys to select RPM Groups and press the Enter key. The screen in Figure 2-65 on page 74 appears.


```

YaST @ LINUXFW                                     Press F1 for Help

+RPM Groups-----+formation-]          [Etc.-]
|--Amusements
|---Development
|++Languages
|+--C and C++
|+--Fortran
|+--Java
|+--Other
|+--Perl
|+--Python
|+--Scheme
|+--Tcl
|++Libraries
|+--Sources
|++Tools
|--Documentation
|--Hardware
|--Productivity
|--Programming
+-----+-----+-----+
|Name      |Avail. Vers.|Inst. Vers.|Summary
|python    |2.3.3       |            |Python Interpreter
|python-64bit|9           |            |Python Interpreter
|python-demo|2.3.3      |            |Python Demonstration Scripts
|python-devel|2.3.3     |            |Include Files and Libraries Ma
|python-doc |2.3.3     |            |Additional Package Documentati
|python-doc-pdf|2.3.3    |            |Python PDF Documentation
|python-idle|2.3.3     |            |Integrated development envirom
+-----+-----+-----+
Filter: Python                                     Required Disk Space: 56.2 MB
+-----+-----+-----+
|python - Python Interpreter
|Version: 2.3.3-88.1 Size: 15.1 MB Media No.: 2
|License: Python
|Package Group: Development/Languages/Python
|Provides: _bsddb.so, _csv.so, _hotshot.so, _locale.so, _random.so,
|_socket.so, _ssl.so, _testcapi.so, _weakref.so, array.so, audioop.so,
+-----+-----+-----+
[Help-]    [Search]    [Disk Usage]          [Cancel]    [Accept]

```

Figure 2-67 YaST - Languages expanded

- m. Use the arrow keys to select Python and press the Enter key. The screen in Figure 2-68 will appear.

```

YaST @ LINUXFW                                     Press F1 for Help

[Filter-]          [Actions-]          [Information-]          [Etc.-]
+-----+-----+-----+
|Name      |Avail. Vers.|Inst. Vers.|Summary
|python    |2.3.3       |            |Python Interpreter
|python-64bit|9           |            |Python Interpreter
|python-demo|2.3.3      |            |Python Demonstration Scripts
|python-devel|2.3.3     |            |Include Files and Libraries Ma
|python-doc |2.3.3     |            |Additional Package Documentati
|python-doc-pdf|2.3.3    |            |Python PDF Documentation
|python-idle|2.3.3     |            |Integrated development envirom
+-----+-----+-----+
Filter: Python                                     Required Disk Space: 56.2 MB
+-----+-----+-----+
|python - Python Interpreter
|Version: 2.3.3-88.1 Size: 15.1 MB Media No.: 2
|License: Python
|Package Group: Development/Languages/Python
|Provides: _bsddb.so, _csv.so, _hotshot.so, _locale.so, _random.so,
|_socket.so, _ssl.so, _testcapi.so, _weakref.so, array.so, audioop.so,
+-----+-----+-----+
[Help-]    [Search]    [Disk Usage]          [Cancel]    [Accept]

```

Figure 2-68 YaST - Available python RPMs

- n. Use the arrow keys to select python from the list, and press the Enter key.
- o. Use the arrow keys to select python-64bit from the list, and press the Enter key.

- p. Use the arrow keys to select python-devel from the list and press the Enter key. Your screen should look similar to Figure 2-69, with plus symbols (+) beside each selected RPM.

```

YaST @ LINUXFW                                     Press F1 for Help

[Filter-]      [Actions-]      [Information-]      [Etc.-]
+-----+-----+-----+-----+
|  |Name      |Avail. Vers.|Inst. Vers.|Summary
| + |python    |2.3.3      |           |Python Interpreter
| + |python-64bit|9          |           |Python Interpreter
|   |python-demo|2.3.3      |           |Python Demonstration Scripts
| + |python-devel|2.3.3      |           |Include Files and Libraries Ma
|   |python-doc  |2.3.3      |           |Additional Package Documentati
|   |python-doc-pdf|2.3.3     |           |Python PDF Documentation
|   |python-idle |2.3.3      |           |Integrated development environm
+-----+-----+-----+-----+
Filter: Python                                     Required Disk Space: 60.1 MB
+-----+-----+-----+-----+
|python-64bit - Python Interpreter                |
|Version: 9-200407011606 Size: 3.8 MB Media No.: 3|
|License: GPL, LGPL                              |
|Package Group: Development/Languages/Python     |
|Provides: _bsddb.so()(64bit), _csv.so()(64bit),  |
|_locale.so()(64bit), _random.so()(64bit), _socket.so()(64bit),|
+-----+-----+-----+-----+
[Help-]      [Search]      [Disk Usage]      [Cancel]      [Accept]

```

Figure 2-69 YaST - Selected python RPMs

- q. Press Alt+F to load the Filter menu again.
- r. Use the arrow keys to select RPM Groups and press the Enter key.
- s. Use the arrow keys to select Productivity and press Shift + the equal sign key (=) to expand the menu, as shown in Figure 2-70 on page 77.

```

YaST @ LINUXST                                     Press F1 for Help

+RPM Groups-----+formation-]          [Etc.-]
|+-Amusements
|+-Development
|+-Documentation
|+-Hardware
|---Productivity
|++Archiving
|++Clustering
|++Databases
|++Editors
|+-File utilities
|++Graphics
|++Multimedia
|++Networking
|++Office
|---Other
|++Publishing
|++Scientific
|+-Security
+-----+-----+-----+
|s.|Summary
|Recovers damaged tar-archives
|ACE archive extractor
|Pack Program
|A program to extract, test, and v-
|A program to unpack compressed fi
|File compression program
|Pack Program
+-----+-----+-----+
|quired Disk Space: 0 B
+-----+-----+-----+
|iles
|- No.: 1
|BSD
|mpression
+-----+-----+-----+
|u>
+-----+-----+-----+
|                                     [Cancel]   [Accept]

```

Figure 2-70 YaST - Productivity expanded

- t. Use the arrow keys to select Archiving and press Shift + the equal sign key (=) to expand the menu, as shown in Figure 2-71.

```

YaST @ LINUXST                                     Press F1 for Help

+RPM Groups-----+formation-]          [Etc.-]
|+-Amusements
|+-Development
|+-Documentation
|+-Hardware
|---Productivity
|+-Archiving
| |+-Backup
| |+-Compression
|++Clustering
|++Databases
|++Editors
|+-File utilities
|++Graphics
|++Multimedia
|++Networking
|++Office
|---Other
|++Publishing
+-----+-----+-----+
|s.|Summary
|Recovers damaged tar-archives
|ACE archive extractor
|Pack Program
|A program to extract, test, and v-
|A program to unpack compressed fi
|File compression program
|Pack Program
+-----+-----+-----+
|quired Disk Space: 0 B
+-----+-----+-----+
|iles
|- No.: 1
|BSD
|mpression
+-----+-----+-----+
|u>
+-----+-----+-----+
|                                     [Cancel]   [Accept]

```

Figure 2-71 YaST - Archiving expanded

- u. Use the arrow keys to select Compression and press the Enter key. The screen in Figure 2-72 on page 78 appears.

```

YaST @ LINUXST                                     Press F1 for Help
[Filter-]      [Actions-]      [Information-]      [Etc.-]
+-----+-----+-----+-----+
| Name      | Avail. Vers. | Inst. Vers. | Summary
| tarfix    | 1.0          |              | Recovers damaged tar-archives
| unace     | 1.2b        |              | ACE archive extractor
| unarj     | 2.65        |              | Pack Program
| unrar     | 3.3.6       |              | A program to extract, test, and v-
| unzip     | 5.50        |              | A program to unpack compressed fi
| zip       | 2.3         |              | File compression program
| zoo       | 2.10        |              | Pack Program
+-----+-----+-----+-----+
Filter: Compression                               Required Disk Space: 0 B
+-----+-----+-----+-----+
| unzip - A program to unpack compressed files
| Version: 5.50-345.1 Size: 649.0 kB Media No.: 1
| License: Other License(s), see package, BSD
| Package Group: Productivity/Archiving/Compression
| Provides: crunzip
| Authors: Info-ZIP <zip-bugs@lists.wku.edu>
+-----+-----+-----+-----+
[Help-]      [Search]      [Disk Usage]      [Cancel]      [Accept]

```

Figure 2-72 YaST - Available compression RPMs

- v. Use the arrow keys to select unzip from the list, and press the Enter key. Your screen should look similar to Figure 2-73, with plus symbols beside each selected RPM.

```

YaST @ LINUXST                                     Press F1 for Help
[Filter-]      [Actions-]      [Information-]      [Etc.-]
+-----+-----+-----+-----+
| Name      | Avail. Vers. | Inst. Vers. | Summary
| tarfix    | 1.0          |              | Recovers damaged tar-archives
| unace     | 1.2b        |              | ACE archive extractor
| unarj     | 2.65        |              | Pack Program
| unrar     | 3.3.6       |              | A program to extract, test, and v-
| + unzip   | 5.50        |              | A program to unpack compressed fi
| zip       | 2.3         |              | File compression program
| zoo       | 2.10        |              | Pack Program
+-----+-----+-----+-----+
Filter: Compression                               Required Disk Space: 737.0 kB
+-----+-----+-----+-----+
| unzip - A program to unpack compressed files
| Version: 5.50-345.1 Size: 649.0 kB Media No.: 1
| License: Other License(s), see package, BSD
| Package Group: Productivity/Archiving/Compression
| Provides: crunzip
| Authors: Info-ZIP <zip-bugs@lists.wku.edu>
+-----+-----+-----+-----+
[Help-]      [Search]      [Disk Usage]      [Cancel]      [Accept]

```

Figure 2-73 YaST - Selected compression RPMs

- y. When you are prompted to swap CDs, do so and press the Enter key. Remember that you may have to press the Enter key twice before the CD will be read. Installation of the GCC tools requires SLES9 CDs two through four.
 - z. When installation of the GCC tools has been completed, the main YaST screen will appear again. Tab to the Quit button and press the Enter key.
33. You will be back at the Linux command shell. Your Linux operating system is now setup and ready to be secured.

2.3.2 Installing RHEL4

Note: The following steps give instructions for installing the RHEL4 operating system on one of the LPARs you have just created. You must repeat these steps to install RHEL4 on both of the new LPARs. On the first iteration of the steps, use the values in column 2 on the planning worksheet to set up the firewall partition. On the second iteration of the steps, use the values from column 3 to set up the security tools partition.

To install Linux on your LPAR, follow these steps:

1. Insert the first disk from your Linux distribution into the CD-ROM drive on your eServer i5.
2. Open the virtual console connection:
 - a. Open up the PuTTY telnet client.
 - b. Enter your i5/OS server name as the Host Name (or IP address).
 - c. Set the protocol to Telnet.
 - d. Enter 2301 as the Port. Your settings should look similar to Figure 2-34 on page 51. Click the **Open** button to open the connection.




Figure 2-76 PuTTY settings

- e. You will be presented with a screen similar to the one shown in Figure 2-35 on page 51. Enter the number that corresponds with the appropriate NWSD name from row D on the planning worksheet, and press the Enter key.

```
OS/400 Guest Partition Consoles
2: Linux_ST(V1-C2/V2-C0)
3: Linux_FW(V1-C4/V3-C0)

Enter the console partition number:
```

Figure 2-77 OS/400 Guest Partition Consoles

- f. You will be asked to enter your service tools user ID. Enter it now, and press the Enter key.
- g. Enter your service tools user ID password, and press the Enter key. You will see the screen shown in Figure 2-36 on page 52. Leave this PuTTY terminal open.

```
Linux_FW: Enter OS/400 service tools userid:
linuxusr
Linux_FW: Enter OS/400 service tools password:
Linux_FW: Console connecting...
Linux_FW: Console connected.
```

Figure 2-78 Console connected

3. Set the NWSD boot parameters:
 - a. Open a 5250 command line to your eServer i5.
 - b. Enter the WRKNWSD command, and press the Enter key.
 - c. Place a 2 beside LINFWSD and press the Enter key.
 - d. Page down twice until you see the screen shown in Figure 2-37 on page 52.

```

Change Network Server Desc (CHGNWSD)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . *SYS
+ for more values
Restricted device resources . . *NONE      Name, *SAME, *NONE, *ALL...
+ for more values
Synchronize date and time . . . *NO      *SAME, *TYPE, *YES, *NO
IPL source . . . . . *STMF      *SAME, *NWSSTG, *PANEL...
IPL stream file . . . . . *NONE

IPL parameters . . . . . *NONE

More...
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 2-79 Change NWSD

- e. Change the value of IPL stream file to `‘/Q0PT/RED_HAT/IMAGES/PSERIES/NETBOOT.IMG’`.
- f. Press the Enter key to save the changes. You will see the message Description for network server changed printed at the bottom of the screen.
- g. Press the F3 key.
4. Vary on the NWSD:
 - a. Type `WRKCFGSTS (*NWS)` at the 5250 command line, and press the Enter key.
 - b. You will be shown a list of NWSDs currently on your system, along with their status, as shown in Figure 2-38 on page 53.

```

Work with Configuration Status                                RCHAS10
                                                           04/25/05 10:12:06
Position to . . . . . Starting characters
Type options, press Enter.
 1=Vary on  2=Vary off  5=Work with job  8=Work with description
 9=Display mode status 13=Work with APPN status...

Opt  Description      Status      -----Job-----
    LINFWSO          VARIED OFF
    LINSTSD          VARIED OFF

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

```

Figure 2-80 Work with Configuration Status screen

- c. Type a 1 beside the appropriate NWSD name from row D on the planning worksheet, and press Enter.
 - d. The status of the NWSD will change to VARY ON PENDING, and the message Vary on completed for network server will be printed at the bottom of the screen.
 - e. If you refresh the screen by pressing the F5 key, the status should read ACTIVE. Do not end this 5250 session; leave it open.
5. Back in the PuTTY terminal that you opened earlier, you should now begin to see output from the LPAR boot sequence. The first boot messages you see should look like the screen in Figure 2-39 on page 54.

7. The welcome screen shown in Figure 2-83 is shown. Tab to the OK button and press the Enter key.




Figure 2-83 Anaconda - Welcome

8. The language selection screen appears, as shown in Figure 2-84 on page 86. Use the arrow keys to select your language. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Language Selection +-----+
|
| What language would you like to use
| during the installation process?
|
| Catalan #
| Chinese(Simplified) #
| Chinese(Traditional) #
| Croatian #
| Czech #
| Danish #
| Dutch #
| English #
|
| +----+ +-----+
| | OK | | Back |
| +----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-84 Anaconda - Language selection

9. Anaconda will search for existing RHEL installations, as shown in Figure 2-85. Be patient while this process finishes.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Searching +-----+
|
| Searching for Red Hat Enterprise Linux AS installations...
|
| 0%
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-85 Anaconda - Searching for existing RHEL installations

10. Autopartition the system:

- a. Tab to the Autopartition option on the screen shown in Figure 2-86 and press the Enter key.




Figure 2-86 Anaconda - Partitioning options

- b. Accept the default options on the screen shown in Figure 2-87 on page 88. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.
+-----+ Automatic Partitioning +-----+
|
| Before automatic partitioning can be set up by
| the installation program, you must choose how to
| use the space on your hard drives.
|
| Remove all Linux partitions on this system
| Remove all partitions on this system
| Keep all partitions and use existing free space
|
| Which drive(s) do you want to use for this installation?
| [*] sda #
|      #
|      #
|
| +----+ +-----+
| | OK | | Back |
| +----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-87 Anaconda - Choose how to use the available space

- c. Confirm to Anaconda that you do want to erase all data on the partition by tabbing to the Yes button on the screen shown in Figure 2-88 and pressing the Enter button.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.
+-----+ Automatic Partitioning +-----+
|
| Before automatic partitioning can be set up by
| the installation program, you must choose how to
| use the space on your hard drives.
|
| You have chosen to remove all partitions
| (ALL DATA) on the following drives:
| /dev/sda
|
| Are you sure you want to do this?
|
| +----+ +-----+
| | No | | Yes |
| +----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-88 Anaconda - Confirm hard drive erase

- d. Accept the default partitioning information that Anaconda creates. It should look similar to Figure 2-89. Tab to the OK button and press the Enter key.




Figure 2-89 Anaconda - Suggested partitions

11. Configure networking:

- a. Anaconda detects the first network in the system and displays the associated property page, as shown in Figure 2-90 on page 90. Tab to Configure using DHCP and press the Spacebar to deselect the option.

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

```
+-----+ Network Configuration for eth0 +-----+
|
| Network Device: eth0
|
| [*] Configure using DHCP
| [*] Activate on boot
|
| IP Address      _____
| Netmask        _____
|
|      +-----+      +-----+
|      | OK  |      | Back |
|      +-----+      +-----+
|
+-----+
```

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

Figure 2-90 Anaconda - Network setup

- b. Use the Tab key to navigate through the fields. Fill in the appropriate value for the IP Address and Netmask from the planning worksheet. If you are setting up the LINUXFW partition, you should use cells CB2 and CC2. If you are setting up the LINUXST partition, you should use cells AB3 and AC3.
- c. Tab to the OK button and press the Enter key.
If you are setting up the LINUXFW partition, you must now set up the remaining network cards in the same fashion in which the first card was just set up. From the planning worksheet, use cells BB2 and BC2 for eth1 and cells AB2 and AC2 for eth2.
- d. You will be presented with the Miscellaneous Network Settings page shown in Figure 2-91 on page 91. Enter the appropriate value for the Gateway. For the LINUXFW partition, this comes from cell K2 on the planning worksheet. For the LINUXST partition, the gateway is the IP address of the LINUXFW partition, cell AB2 on the planning worksheet.

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

```
++ Miscellaneous Network Settings ++
|
| Gateway: _____
| Primary DNS: _____
| Secondary DNS: _____
| Tertiary DNS: _____
|
| +----+      +-----+
| | OK |      | Back |
| +----+      +-----+
|
+-----+
```

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

Figure 2-91 Anaconda - Miscellaneous Network Settings

- e. Enter the Primary DNS and Secondary DNS from rows L and M on the planning worksheet.
- f. Tab to the OK button and press the Enter key.
- g. Enter in the host name for the partition on the screen shown in Figure 2-92 on page 92. The host name should come from row N on the planning worksheet. Tab to the OK button and press the Enter key.

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

```
+-----+ Hostname Configuration +-----+
|
| If your system is part of a larger network where hostnames are
| assigned by DHCP, select automatically via DHCP. Otherwise,
| select manually and enter in a hostname for your system. If you
| do not, your system will be known as 'localhost.'
|
| ( ) automatically via DHCP
| (*) manually
|
| +----+
| | OK |
| +----+
|
| +-----+
| | Back |
| +-----+
|
+-----+
```

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

Figure 2-92 Anaconda - Hostname settings

- h. The firewall settings screen is displayed, as shown in Figure 2-93. Leave the Enable firewall option selected. Tab to the OK button and press the Enter key.

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

```
+-----+ Firewall +-----+
|
| A firewall can help prevent unauthorized access
| to your computer from the outside world. Would
| you like to enable a firewall?
|
| (*) Enable firewall ( ) No firewall
|
| +----+
| | OK |
| +----+
|
| +-----+
| | Customize |
| +-----+
|
| +----+
| | Back |
| +----+
|
+-----+
```

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

Figure 2-93 Anaconda - Firewall settings

12. Security Enhanced Linux (SELinux) settings are displayed, as shown in Figure 2-94. Leave the Active setting selected, tab to the OK button, and press the Enter key.




Figure 2-94 Anaconda - SELinux settings

13. The Language Support screen shown in Figure 2-95 on page 94 is displayed. If you require additional languages, select them from the list. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Language Support +-----+
|
| Choose additional languages that you would
| like to use on this system:
|
|   [ ] English (New Zealand)           #
|   [ ] English (Philippines)          #
|   [ ] English (Singapore)            #
|   [ ] English (South Africa)         #
|   [*] English (USA)                  #
|   [ ] English (Zimbabwe)             #
|   [ ] Estonian                       #
|   [ ] Faroese (Faroe Islands)        #
|
| +----+ +-----+ +-----+ +-----+
| | OK | | Select All | | Reset | | Back |
| +----+ +-----+ +-----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-95 Anaconda - Additional language support

14. Select your time zone from the screen shown in Figure 2-96. Use the arrow keys to scroll through the list. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Time Zone Selection +-----+
|
| What time zone are you located in?
|
| [ ] System clock uses UTC
|
| America/Montevideo                   #
| America/Montreal                     #
| America/Montserrat                   #
| America/Nassau                       #
| America/New_York                     #
|
| +----+ +-----+
| | OK | | Back |
| +----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-96 Anaconda - Select time zone

15. The root password selection screen appears, as shown in Figure 2-97. Enter in your root password from row J on the planning worksheet twice. Tab to the OK button and press the Enter key.




Figure 2-97 Anaconda - Setting the root password

16. Select the required packages for installation:

- a. Tab to the Customize software selection option on the screen shown in FIGGY. Press the Spacebar to select it, tab to the OK button, and press the Enter key.

```
Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Package Defaults +-----+
|
| The default installation environment includes our
| recommended package selection. After installation,
| additional software can be added or removed using
| the 'system-config-packages' tool.
|
| However Red Hat Enterprise Linux AS ships with many
| more applications, and you may customize the
| selection of software installed if you want.
|
| [ ] Customize software selection
|
| +-----+           +-----+
| | OK |             | Back |
| +-----+           +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen
```

Figure 2-98 Anaconda - Package options

- b. You will be presented with the package selection list shown in Figure 2-99 on page 97. Use the arrow keys and the Spacebar to deselect the following options from the list:
 - X Window System
 - GNOME Desktop Environment
 - Web Server
 - Windows File Server
 - Printing Support
- c. Use the arrow keys and the Spacebar to select the following option:
 - Development Tools
- d. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Package Group Selection +-----+
|
| Total install size: 1,301M
|
| [ ] X Window System #
| [ ] GNOME Desktop Environment #
| [ ] KDE (K Desktop Environment) #
| [ ] Editors #
| [ ] Engineering and Scientific #
| [ ] Graphical Internet #
| [*] Text-based Internet #
| [ ] Office/Productivity #
|
|          +----+          +-----+
|          | OK |          | Back |
|          +----+          +-----+
|
+-----+

<Space>,<+>,<-> selection | <F2> Group Details | <F12> next screen

```

Figure 2-99 Anaconda - Package selection list

17. You will be informed that installation is about to begin in the screen shown in Figure 2-100. Tab to the OK button and press the Enter key.

```

Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Installation to begin +-----+
|
| A complete log of your installation will
| be in /root/install.log after rebooting
| your system. You may want to keep this
| file for later reference.
|
|          +----+          +-----+
|          | OK |          | Back |
|          +----+          +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 2-100 Anaconda - Installation to begin

18. You will be informed which CDs are required for installation, as shown in Figure 2-101. Ensure that you have all the required media, tab to the Continue button, and press the Enter key.

```
Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Required Install Media +-----+
|
| The software you have selected to install will
| require the following CDs:
|
|     Red Hat Enterprise Linux AS 4 CD #1
|     Red Hat Enterprise Linux AS 4 CD #2
|     Red Hat Enterprise Linux AS 4 CD #3
|
| Please have these ready before proceeding with the
| installation. If you need to abort the installation
| and reboot please select "Reboot".
|
|     +-----+           +-----+
|     | Reboot |           | Continue |
|     +-----+           +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen
```

Figure 2-101 Anaconda - Required media

19. The install image will be transferred to the hard drive, as shown in Figure 2-102 on page 99.




Figure 2-102 Anaconda - Transferring install image

20. The Red Hat Package Manager (RPM) transaction will be setup, as shown in Figure 2-103.




Figure 2-103 Anaconda - RPM transaction setup

21. Installation will proceed. When you are asked to swap CDs, as shown in Figure 2-104, do so, and press the Enter key.

```
Red Hat Enterprise Linux AS (C) 2004 Red Hat, Inc.

+-----+ Package Installation +-----+
|
| Name   : hwdata-0.146.1.EL-1-noarch
| Size   +-----+ Change CDROM +-----+
| Summar |
|        | Please insert Red Hat Enterprise Linux
|        | AS disc 2 to continue.
|        |
|        | +-----+
|        | | OK |
|        | +-----+
| Total  |
| Compl  |
| Remai  |
|
+-----+
|
| <Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen
|
+-----+
```

Figure 2-104 Anaconda - Swap CDs

22. You may receive the error shown in Figure 2-105 on page 101. This is normal. Simply press the Enter key again, and installation will continue.

24. You will be presented with a congratulatory message that informs you of a successful installation, as shown in Figure 2-107. Tab to the Reboot button and press the Enter key.




Figure 2-107 Anaconda - RHEL4 successfully installed

25. The system will start to reboot. Vary off the LPAR:

- a. Back in the 5250 screen that you opened earlier, place a 2 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- b. The status of the NWSD will change to VARY OFF PENDING. It will take several minutes before the NWSD is varied off. Occasionally press the F5 key to refresh the screen. When the status of the NWSD reads VARIED OFF, proceed to the next step.
- c. Press the F3 key to exit the Work with Configuration Status screen.

26. Change the LPAR boot options to boot from the NWSSTG:

- a. At the 5250 command line, enter WRKNWSD and press the Enter key.
- b. Place a 2 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- c. Page down twice so that you have the IPL options on your screen.
- d. Change IPL source to *NWSSTG.
- e. Change IPL stream file to *SAME. Your options should look like Figure 2-108 on page 103.




Figure 2-108 Changed NWSD IPL settings

- f. Press the Enter key to save the changed NWSD settings. The message Description for network server changed will be printed at the bottom of your screen.
 - g. Press the F3 key to exit the Change Network Server Description screen.
27. Change the LPAR profile boot options:
- a. In the main HMC window, expand the appropriate LPAR name from row B of the planning worksheet.
 - b. Right-click the **Default** profile, and select **Properties**.
 - c. The Logical Partition Profile Properties window appears.
 - i. Click the **Settings** tab.
 - ii. Under the Boot Modes heading, select **Normal**, as shown in Figure 2-109 on page 104.




Figure 2-109 Changing boot modes

iii. Click the **OK** button to save the settings.

28. Vary on the NWSD:

- a. At a 5250 command line to your eServer i5, type `WRKCFGSTS (*NWS)` and press the Enter key.
- b. You will be shown a list of NWSDs currently on your system, along with their status. Type a 1 beside the appropriate NWSD name from row D of the planning worksheet, and press the Enter key.
- c. The status of the NWSD will change to `VARY ON PENDING`, and the message `Vary on completed for network server` will be printed at the bottom of the screen.
- d. If you refresh the screen by pressing the F5 key, the status should read `ACTIVE`. Do not end this session; leave it open.

29. Back in your PuTTY terminal window, you can see your Linux operating system booting up. The boot process is not interactive, and requires no input from the user.

30. Log in as root, using the password from row J of the planning worksheet.

31. Add a regular user to the system:

- a. At the Linux shell, type the following and press the Enter key:

```
# adduser USERNAME
```

Where *USERNAME* is replaced with the user name of the user you wish to add to the system.

- b. Change the password of the new user:

```
# passwd USERNAME
```

- Where *USERNAME* is replaced with the user name you entered in step a on page 104.
- c. Enter the password you would like the user to have, and press the Enter key.
 - d. Re-enter the password when prompted, and press the Enter key.
 - e. The user has been created. The output should look similar to Figure 2-110.

```
LINUXFW # adduser arobar
LINUXFW # passwd arobar
Changing password for user arobar.
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
LINUXFW #
```

Figure 2-110 Adding a user to the system

32.RHEL4 has been installed successfully on your system, and is ready to be secured.

Archived



Locking down the Linux firewall partition

This chapter describes how to secure your newly installed Linux system. Securing the server includes hardening the Linux installation, configuring a proper firewall, setting up a network intrusion detection system, setting up root kit detection software, and securing remote network access. Note that this process only needs to be performed on the LINUXFW partition, as the LINUXST partition does not have any direct connections to any network.

Important: These steps will lock down your Linux installation. Carefully follow each direction properly, or you may become locked out of your own system.

3.1 Hardening Linux

Hardening the Linux installation means providing greater security measures at the operating system level than those that are provided by default. Hardening the installation involves removing unnecessary servers, updating basic settings to be more secure in a production environment, patching the kernel with a security patch, and updating the kernel with the grsecurity and Security Enhanced Linux (SELinux) patches. The hardening process will also involve setting up logging so that attackers leave a trail of what was used to compromise your system, if they have managed to crack your security.

Note: The planning worksheet referenced in this chapter is located in 1.3.6, “Planning worksheet” on page 13.

3.1.1 Bastille Linux

Bastille Linux is an interactive program written in perl that asks you questions about system security. After the questioning period, Bastille uses your responses to lock down your system. The Bastille Linux program seeks to educate system administrators about security standards in addition to securing the target system. We recommend that you read each question and the accompanying information before selecting an answer.

Installing Bastille Linux

Bastille is not included with SLES9 or RHEL4. It is OSS, and can be downloaded. Follow these instructions to obtain, install, and use Bastille Linux:

1. Now that networking on your Linux system is functional, we recommend that you use SSH to complete administrative tasks instead of using the virtual console. SSH is a much more secure way to administer your Linux operating system than through an openly viewable console. SSH to the firewall partition using PuTTY:
 - a. Open up the PuTTY client.
 - b. Enter the IP address of the LINUXFW partition from cell CB2 on the planning worksheet in the Host Name (or IP address) field.
 - c. Set the protocol to SSH.
 - d. Click the **Open** button to open the connection.
 - e. Log in to your system as the root user, using the password from cell J2 of the planning worksheet.
2. Create a /etc/tools directory for organizational purposes:

```
~> mkdir /etc/tools
```
3. Switch to the new directory:

```
~> cd /etc/tools
```
4. Download the perl Curses module:

```
~> curl -O http://cpan.org/modules/by-module/Curses/Curses-1.12.tgz
```
5. Extract the perl Curses module:

```
~> tar zxf Curses-1.12.tgz
```
6. Switch to the Curses directory:

```
~> cd Curses-1.12
```

7. Compile and install the perl Curses module:

```
~> perl Makefile.PL; make; make install
```

8. Exit the Curses directory:

```
~> cd ..
```

9. Download the Bastille Linux RPM file:

```
~> curl -O
http://easynews.dl.sourceforge.net/sourceforge/bastille-linux/Bastille-3.0.4-1.0.noarch.
rpm
```

10. Install the Bastille Linux RPM:

```
~> rpm -ivh Bastille-3.0.4-1.0.noarch.rpm
```

11. On RHEL4 systems, you will need to create an empty file called rhel4:

```
~> touch /etc/Bastille/rhel4
```

This is not required on SLES9 systems.

Running Bastille Linux on SLES9

Bastille is run using an ncurses-based text interface. Run Bastille:

1. While you are still in the previously opened SSH session to the firewall, run the Bastille program:

```
~> bastille -c
```

2. Accept the Bastille license, shown in Figure 3-1 on page 110. Press the Enter key until you are given an angle bracket prompt. Type accept and press the Enter key.

Note: Remember that these settings are to be changed on the firewall partition only. Using the Bastille Linux software on the LINUXST partition could result in a system that cannot perform proper mail scanning or filtering.

```
[root@LINUXFW tools]# bastille -c
NOTE:    Using Curses user interface module.
NOTE:    Only displaying questions relevant to the current configuration.

Copyright (C) 1999-2002 Jay Beale
Copyright (C) 1999-2001 Peter Watkins
Copyright (C) 2000 Paul L. Allen
Copyright (C) 2001-2003 Hewlett-Packard Development Company, L.P.
Bastille is free software; you are welcome to redistribute it under
certain conditions.  See the 'COPYING' file in your distribution for terms.

DISCLAIMER.  Use of Bastille can help optimize system security, but does not
guarantee system security.  Information about security obtained through use of
Bastille is provided on an AS-IS basis only and is subject to change without
notice.  Customer acknowledges they are responsible for their system's security.
TO THE EXTENT ALLOWED BY LOCAL LAW, Bastille (SOFTWARE) IS PROVIDED TO YOU
AS IS WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, WHETHER ORAL OR WRITTEN,
EXPRESS OR IMPLIED.  JAY BEALE, THE BASTILLE DEVELOPERS, AND THEIR SUPPLIERS
DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.
Some countries, states and provinces do not allow exclusions of implied
warranties or conditions, so the above exclusion may not apply to you.  You may
have other rights that vary from country to country, state to state, or province
to province.  EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT WILL
JAY BEALE, THE BASTILLE DEVELOPERS, OR THEIR SUBSIDIARIES, AFFILIATES OR
SUPPLIERS BE LIABLE FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL OR OTHER
DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF
THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED
IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED
OF THE POSSIBILITY OF SUCH DAMAGES.  Your use of the Software is entirely at your
own risk.  Should the Software prove defective, you assume the entire cost of all
service, repair or correction.  Some countries, states and provinces do not allow
the exclusion or limitation of liability for incidental or consequential
damages, so the above limitation may not apply to you.

You must accept the terms of this disclaimer to use
Bastille.  Type "accept" (without quotes) within 5
minutes to accept the terms of the above disclaimer
>
```

Figure 3-1 Bastille - License

3. The Bastille user interface appears, as shown in Figure 3-2 on page 111. Tab to Next and press the Enter key.




Figure 3-2 Bastille - User interface

4. Restrict administrative functions, as shown in Figure 3-3 on page 112.

Restricting administrative functions to only the root user forces a would-be attacker to become the root user before using programs that can be potentially damaging to your system. Becoming root is much more difficult than gaining access to a standard user account, so it is prudent to restrict administrative functions. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to set more restrictive permissions on the administration
|utilities? [N]
|In general, the default file permissions set by most vendors are fairly secure
|. To make them more secure, though, you can remove non-root user access to
|some administrator functions.
|
|If you choose this option, you'll be changing the permissions on some common
|system administration utilities so that they're not readable or executable by
|users other than root. These utilities (which include linuxconf, fsck,
|ifconfig, runlevel and portmap) are ones that most users should never have a
|need to access. This option will increase your system security, but there's a
|chance it will inconvenience your users.
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back > < Next > < Explain Less >

```

Figure 3-3 Bastille - Restricting administrative functions

5. Restrict file permission, as shown in Figure 3-4.

By restricting system files to only the root user, you are enforcing a would-be attacker to become the root user to change configuration which can potentially damage your system. Becoming root is much more difficult than gaining access to a standard user account, so it is prudent to restrict file permissions. Tab to the Next option and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|The following questions all pertain to disabling "SUID root" permission for
|particular programs. This permission allows non-root users to run these
|programs, increasing convenience but decreasing security. If a security
|weakness or vulnerability is found in these programs, it can be exploited to
|gain root-level access to your computer through any user account.
|
|If you answer "Yes" and then realize later that you do need SUID permissions
|on a specific program, you can always turn it back on later with chmod u+s <
|file name>.
|
+-----+
|
| < Back > < Next > < Explain Less >

```

Figure 3-4 Bastille - Restricting file permissions

6. Disable SUID for the mount and umount commands, as shown in Figure 3-5.

This forces an attacker to become root before they can mount or unmount any file systems on your server. Use the arrow keys to select Yes and press the Enter key.

Note: Forcing everyone to become the root user before they can use any kind of tool that may damage your system provides you with user accountability. Direct login to your system as the root user will be disabled later on in the Bastille program. This forces a user to log in as a local standard user before becoming root. Any damage done as root can then be traced back to the standard user account that an attacker initially used to log in to the system. With knowledge of this, most users will opt to keep their passwords a secret. In the event that your system is compromised, knowing which user account was used to start with provides you with a starting point for your investigation into the attack.

```
+-----+
|                                     |
|                               Bastille |
|-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for mount/umount?
|Mount and umount are used for mounting (activating) and unmounting (
|deactivating) drives that were not automatically mounted at boot time. This
|can include floppy and CD-ROM drives. Disabling SUID would still allow anyone
|with the root password to mount and unmount drives.
|
|Would you like to disable SUID status for mount/umount?
|
+-----+
|                                     |
|                               +-----+ |
|                               |Yes  | |
|                               |No   | |
|                               +-----+ |
|                                     |
| < Back >   < Next >   < Explain Less > |
+-----+
```

Figure 3-5 Bastille - Disabling SUID for the mount and umount commands

7. Disable SUID for the ping command, as shown in Figure 3-6 on page 114.

This forces an attacker to become root before they can use the ping command, as the ping command uses root privileges to open a raw socket. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for ping? [Y]
|Ping is used for testing network connectivity. Specifically it's for testing
|the ability of the network to get a packet from this machine to another and
|back. The ping program is SUID since only the root user can open a raw socket
|. Since, however, it is often used only by the person responsible for
|networking the host, who normally has root access, we recommend disabling SUID
|status for it.
|
|Would you like to disable SUID status for ping? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back > < Next > < Explain Less >

```

Figure 3-6 Bastille - Disabling SUID for the ping command

8. Disable SUID for the **at** command, as shown in Figure 3-7.

This forces standard users to use the cron system to schedule jobs, so the at scheduling system cannot be exploited. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for at? [Y]
|"at" is used for scheduling an individual task to run at a single later time.
|There have historically been many exploits that take advantage of weaknesses
|in "at". Virtually all of the necessary functionality of "at" can be found in
|cron (and removing cron is not practical) so there is no need to retain
|privileged access for "at".
|
|Would you like to disable SUID status for at? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back > < Next > < Explain Less >

```

Figure 3-7 Bastille - Disabling SUID for the at command

9. Disable the BSD r-tools, as shown in Figure 3-8 on page 115.


```

+-----+
|                               Bastille                               |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Should Bastille disable clear-text r-protocols that use IP-based
|authentication? [Y]
|The BSD r-tools rely on IP-based authentication, which means that you can
|allow anyone with (for instance) root access on 192.168.1.1 have root access
|on 192.168.1.2. Administrators and other users have traditionally found this
|useful, as it lets them connect from one host to another without having to
|retype a password. The .rhosts file contains the names of the accounts and
|machines that are considered to be trusted.
|
|The problem with IP-based authentication, however, is that an intruder can
|craft "spoofed" or faked packets which claim to be from a trusted user on a
|trusted machine. Since the r-tools rely entirely on IP addresses (and remote
|username) for authentication, a spoofed packet will be accepted as real.
+-----+
|                               +-----+
|                               |Yes  |
|                               |No   |
|                               +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-9 Bastille - Disabling the clear-text r-protocols

11. Do not enforce password aging, as shown in Figure 3-10 on page 117.

We recommend that passwords *do* change; however, we recommend that changing the password be part of your administration routine, along with other administrative tasks such as reading the system logs. This option may be safely enabled; however, we recommend that the password aging option remains disabled. You are more likely to lose track of your password if you are forced to change it as opposed to manually changing it. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Would you like to enforce password aging? [Y]
|Your operating system's default behavior, which we would change here, is to
|disable an account when the password hasn't changed in 99,999 days. This
|interval is too long to be useful. We can set the default to 60 days. At
|some point before the 60 days have passed, the system will ask the user to
|change his or her password. At the end of the 60 days, if the password has
|not been changed, the account will be temporarily disabled. We'll make sure
|this warning period is at least 5 days long. We would make this change in /
|etc/login.defs.
|
|Would you like to enforce password aging? [Y]
|
+-----+
|          +-----+
|          | Yes   |
|          | No   |
|          +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-10 Bastille - Leaving the password aging settings

12. Disallow root login on the system consoles, as shown in Figure 3-11.

This forces an attacker to know the password for a standard user account as well as the root account before they can damage your system. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Should we disallow root login on tty's 1-6? [N]
|You can restrict which tty's root can login on. Some sites choose to restrict
|root logins, so that an admin must login with an ordinary user account and
|then use su to become root.
|
|This can stop an attacker who has only been able to steal the root password
|from logging in directly. He has to steal a second account's password to make
|use of the root password via the ttys.
|
|Should we disallow root login on tty's 1-6? [N]
|
+-----+
|          +-----+
|          | Yes   |
|          | No   |
|          +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-11 Bastille - Restricting root login on system consoles


```

+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Would you like to reduce the LILO delay time to zero? [N]
|We can further protect the system by taking away the attacker's chance to type
|anything at the LILO prompt. This is not dependent on the previous option,
|nor is it exclusive of it. If you chose the previous option, this will make
|your configuration even tighter, as some machines will allow an attacker to
|place keystrokes into the keyboard buffer before he or she reaches the LILO
|prompt.
|
|Would you like to reduce the LILO delay time to zero? [N]
|
+-----+
|
| +-----+
| | Yes |
| | No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-13 Bastille - Leaving the LILO delay-time options

15. Inform Bastille that your boot device is your hard drive, as shown in Figure 3-14. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Do you ever boot Linux from the hard drive? [Y]
|If you selected "yes" on either of the previous options (password-protecting
|the LILO prompt or reducing its delay to zero), then you need to now write the
|changes to your LILO configuration.
|
|Do you boot from your hard drive? That is, is LILO installed on your hard
|drive?
|
|Do you ever boot Linux from the hard drive? [Y]
|
+-----+
|
| +-----+
| | Yes |
| | No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-14 Bastille - Setting your hard drive as your boot device

16. Do not write the LILO changes to a floppy, as shown in Figure 3-15 on page 120.

There were no changes made to LILO, and as such, no changes need to be saved to a floppy disk. Use the arrow keys to select No and press the Enter key.

```
+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Would you like to write the LILO changes to a boot floppy? [N]
|If you have a Linux boot floppy, either for normal booting or for emergency
|use, you should also write these LILO changes to that floppy.  If you do not
|already have a customized Linux boot floppy, or if you did not choose to make
|any changes to your LILO configuration, you should answer "no" here.
|
|Would you like to write the LILO changes to a boot floppy? [N]
|
+-----+
|
|          +-----+
|          | Yes   |
|          | No   |
|          +-----+
|
| < Back >   < Next >   < Explain Less >
|
```

Figure 3-15 Bastille - Leaving the LILO floppy disk options

17. Disable Ctrl+Alt+Delete rebooting, as shown in Figure 3-16 on page 121.

Note: The Bastille program states that unless the powerline, power switch, and case of the server can be physically protected, disabling the ability to reboot via the Ctrl+Alt+Delete sequence is unnecessary. However, it is most likely that an attacker will not have physical access to the server. Disabling the Ctrl+Alt+Delete sequence reduces the risk of an attacker being able to reset the server.

There is a good chance that an attacker would be using a virtual console to attack your system, in which case disabling the Ctrl+Alt+Delete reboot is a prudent security measure. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Would you like to disable CTRL-ALT-DELETE rebooting? [N]
|Disabling CTRL-ALT-DELETE rebooting is designed to prevent an attacker with
|access to the machine's keyboard from being able to reboot the machine. A
|reboot done in this manner should not damage the file system, as it shuts the
|machine down cleanly, writing out all pending data in the disk cache to disk
|first. Even with this functionality disabled, however, an attacker could just
|power cycle machine or pull the power cord.
|
|Unless the power line, switch and case of the machine can be physically
|protected, this precaution is wholly unnecessary. Given the fact that the
|attacker _can_ reboot the machine, would you prefer that s/he do it in a way
|potentially damages the file system? Think carefully here, as maintaining the
|integrity of the machine's file system may be secondary to the goal of keeping|
+-----+
|
|          +-----+
|          |Yes  |
|          |No  |
|          +-----+
|
| < Back >  < Next >  < Explain Less >

```

Figure 3-16 Bastille - Disabling Ctrl+Alt+Delete rebooting

18. Default-deny TCP wrappers and xinetd, as shown in Figure 3-17.

The firewall partition should not be using or hosting any network services, so it is safe to deny TCP wrappers and xinetd by default. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+SecureInetd.pm Module 5 of 0-----+
|Q: Would you like to set a default-deny on TCP Wrappers and xinetd? [N]
|Not recommended for most users:
|
|Many network services can be configured to restrict access to certain network
|addresses (and in the case of 'xinetd' services in Linux-Mandrake 8.0 and Red
|Hat 7.x, other criteria as well). For services running under the older 'inetd
|' super-server (found in older versions of Linux-Mandrake and Red Hat, and
|current versions of some other distributions), some standalone services like
|OpenSSH, and --unless otherwise configured-- services running under Red Hat's
|xinetd super-server, you can configure restrictions based on network address
|in /etc/hosts.allow. The services using inetd or xinetd typically include
|telnet, ftp, pop, imap, finger, and a number of other services.
|
+-----+
|
|          +-----+
|          |Yes  |
|          |No  |
|          +-----+
|
| < Back >  < Next >  < Explain Less >

```

Figure 3-17 Bastille - Default-denying TCP wrappers and xinetd

19. Display an authorized use message at login, as shown in Figure 3-18.

This protection measure simply makes it easier for your organization to prosecute an attacker should they damage your system. By providing all users with an authorized use message at login, you are cutting off the argument that an attacker was unaware of what system they were using. Use the arrow keys to select Yes and press the Enter key.

You are informed that the authorized use banner was created in /etc/issue. Press the Tab key to continue.

```
+-----+
|                                     |
|                               Bastille |
+-----+
+SecureInetd.pm Module 5 of 0-----+
|Q: Would you like to display "Authorized Use" messages at log-in time? [Y]
|At this point you can create "Authorized Use Only" messages for your site.
|These may be very helpful in prosecuting system crackers you may catch trying
|to break into your system. Bastille can make default messages which you may
|then later edit. This is sort of like an "anti-welcome mat" for your computer
|.
|
|Would you like to display "Authorized Use" messages at log-in time? [Y]
|
+-----+
|                                     |
|                               +-----+ |
|                               | Yes | |
|                               | No  | |
|                               +-----+ |
|                                     |
| < Back >   < Next >   < Explain Less > |
+-----+
```

Figure 3-18 Bastille - Creating an authorized use message

20. Customize the authorized use message by typing the name or e-mail address of the system administrator, as shown in Figure 3-19 on page 123. You may instead type the name of your organization.

We recommend that you include some method of communication to the system administrator, as it cuts off the argument that an attacker was unable to ask for permission to use the system, because there was no way to contact the system administrator. Use Tab to select the Next button and press the Enter key.

```
+-----+
| Bastille |
+-----+
+SecureInetd.pm Module 5 of 0-----+
|Q: Who is responsible for granting authorization to use this machine?
|Bastille will start to make the banner more specific by telling the user who
|is responsible for this machine. This will state explicitly from whom the
|user needs to obtain authorization to use this machine. Please type in the
|name of the company, person, or other organization who owns or is responsible
|for this machine.
|
|Who is responsible for granting authorization to use this machine?
|
+-----+
| Answer: Alex Robar |
+-----+
| < Back > < Next > < Explain Less > |
+-----+
```

Figure 3-19 Bastille - Customizing the authorized use message

21. Limit system resources, as shown in Figure 3-20 on page 124.

This partition will only be used to protect your network and detect network intrusions. As such, it will never require any user to be running more than 150 simultaneous processes. Use the arrow keys to select Yes and press the Enter key.

You will be informed that the system resource limits have been set in the file /etc/security/limits.conf. Press the Tab key to continue.

```

+-----+
| Bastille |
+-----+
+ConfigureMiscPAM.pm Module 7 of 0-----+
|Q: Would you like to put limits on system resource usage? [N]
|Denial of Service attacks are often very difficult to defend against, since
|they don't require access of any kind to the target machine. Since several
|major daemons, including the web, name, and FTP servers, may run as a
|particular user, you can limit the effectiveness of many Denial of Service
|attacks by modifying /etc/security/limits.conf. If you restrict the resources
|available in this manner, you can effectively cripple most Denial of Service
|attacks.
|
|If you choose this option, you'll be setting the following initial limits on
|resource usage:
|
| - The number of allowed core files will be set to zero. Core files
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-20 Bastille - Limiting system resources

22. Do not restrict console login privileges to a small group of users, as shown in Figure 3-21.

It is feasible for this to be defined; however, the setup for the firewall partition should only include one standard user in addition to the root user. Adding additional users is a security risk, and as such, there should be no need to define which users can log in at a console. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+ConfigureMiscPAM.pm Module 7 of 0-----+
|Q: Should we restrict console access to a small group of user accounts? [N]
|Under some distributions, users logged in at the console have some special
|access rights (like the ability to mount the CD-ROM drive). You can disable
|this special access entirely, but a more flexible option is to restrict
|console access to a small group of trusted user accounts.
|
|Should we restrict console access to a small group of user accounts? [N]
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-21 Bastille - Leaving console group login privileges

23. Enable process accounting, as shown in Figure 3-22.

Should your system be attacked, having process accounting enabled will allow you to view log files that detail what the attacker did to your system, in addition to providing you with insight as to how they cracked your security measures, allowing you to patch the security hole. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
| Bastille |
+-----+
+Logging.pm Module 8 of 0-----+
|Q: Would you like to set up process accounting? [N]
|Linux has the ability to log which commands are run when and by whom. This is
|extremely useful in trying to reconstruct what a potential cracker actually
|ran. The drawbacks are that the logs get large quickly (a log rotate module
|is included to offset this), the parameters to commands are not recorded, and
|, like all log files, the accounting log is removable if the attacker has root
|.
|As this is rather disk and CPU intensive, please choose NO unless you have
|carefully considered this option.
|Would you like to set up process accounting? [N]
+-----+
|
| +-----+
| | Yes |
| | No  |
| |     |
| +-----+
|
| < Back >    < Next >    < Explain Less >
```

Figure 3-22 Bastille - Enabling process accounting

24. Activate the Linux Auditing Subsystem (LAuS), as shown in Figure 3-23 on page 126.

Performing security audits on your system is a good practice to employ. The LAuS system performs this audit automatically, logging all security-related process calls to the kernel. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+Logging.pm Module 8 of 0-----+
|Q: May we activate LAuS?
|The Linux Auditing Subsystem, or LAuS, provides a central security event
|monitoring technology. It logs security-relevant kernel subroutine calls, or
|syscalls, including the parameters the syscalls are called with and the
|success or failure-related return code. The relevant system daemon is auditd.
|
|May we activate LAuS?
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-23 Bastille - Activating LAuS

25. Bastille informs you of how it will approach securing system daemons. Read the explanation, tab to the Next button, and press the Enter key to continue, as shown in Figure 3-24.

```

+-----+
| Bastille |
+-----+
+MiscellaneousDaemons.pm Module 9 of 0-----+
|To make the operating system more secure, we try to deactivate all system
|daemons, especially those running at a high/unlimited level of privilege.
|Each active system daemon serves as a potential point of break-in, which might
|allow an attacker illegitimate access to your system. An attacker can use
|these system daemons to gain access if they are later found to have a bug or
|security vulnerability.
|
|We practice a minimalist principle here: minimize the number of privileged
|system daemons and you can decrease your chances of being a victim should one
|of the standard daemons be found later to have a vulnerability. This section
|will require careful attention, but if you have doubts, you should be able to
|safely select the default value in most cases.
|
+-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-24 Bastille - Securing system daemons

26. Bastille shows you a recommendation for changes to the Apache Web server configuration, as seen in Figure 3-25.

As the firewall partition will not be running any Web server, this can be safely ignored. Tab to the Next button and press the Enter key to continue.

```
+-----+
|                                     Bastille                                     |
+-----+
+Apache.pm Module 12 of 0-----+
| There are a few other changes that we recommend you make to the web server's  |
| configuration. There are very few intrinsic security flaws in the Apache web  |
| server, but there are two important ones:                                     |
|                                                                              |
|   As with all web servers, it is generally required to send and receive     |
|   information to and from anyone on the internet.                           |
|                                                                              |
|   In many environments, the people telling the server how to behave are not   |
|   knowledgeable system administrators by trade. Before you discount this fact|
|   , take account of the wide proliferation of configurations under which any  |
|   user on the system can instruct the server to execute arbitrary code for   |
|   anyone who comes to the site, via CGI scripts.                            |
+-----+

< Back >    < Next >    < Explain Less >
```

Figure 3-25 Bastille - Leaving Apache server settings

27. Do not enable TMPDIR scripts, as shown in Figure 3-26 on page 128.

The security of the /tmp file system will be addressed later in 3.1.3, "Altering insecure defaults" on page 156. Enabling the TMPDIR scripts could interfere with this security measure. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+TMPDIR.pm Module 17 of 0-----+
|Q: Would you like to install TMPDIR/TMP scripts? [N]
|Many programs use the /tmp directory in ways that are dangerous on multi-user
|systems. Many of those programs will use an alternate directory if one is
|specified with the TMPDIR or TMP environment variables. We can install scripts
|that will be run when users log in that safely create suitable temporary
|directories and set the TMPDIR and TMP environment variables. This depends on
|your system supporting /etc/profile.d scripts.
|
|Would you like to install TMPDIR/TMP scripts? [N]
|
+-----+
|          +-----+
|          | Yes   |
|          | No   |
|          +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-26 Bastille - Leaving TMPDIR script settings

28. Do not run the packet filtering script seen in Figure 3-27.

The iptables rules will be set manually, making this script unnecessary. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+Firewall.pm Module 18 of 0-----+
|Q: Would you like to run the packet filtering script? [N]
|Using the packet filtering script, you will be able to do packet filtering/
|modification via the Linux kernel. You can use this to block certain types of
|connections to or from your machine, to turn your machine into a small
|firewall, and to do Network Address Translation (also known as "IP
|masquerading"), which lets several machines share a single IP address.
|
|If you install the packet filtering script, it will create firewalling
|instructions for you. You will be prompted to make various choices (with
|suggested defaults), but you may need to edit it for your particular site and
|WILL need to individually activate it.
|
|This script supports both kernel 2.2 (ipchains) and 2.4 (iptables if available)
|
+-----+
|          +-----+
|          | Yes   |
|          | No   |
|          +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-27 Bastille - Leaving packet filtering script settings


```

+Bastille Credits      (press TAB to go on)-----+
| Jay Beale           - Lead Architect and Original Author |
| HP Bastille Dev Team - Developers - HP-UX Port, Design/Arch. |
| Peter Watkins       - Core Developer: Firewall           |
| Mike Rash           - Developer: PSAD                   |
| Paul Allen          - Developer: User Interface          |
| Javier Fernandez-Sanguino - Developer - Debian Port     |
| Niki Rahimi (IBM)   - Developer - SuSE and TurboLinux Ports |
| Brian Stine         - Developer - Gentoo Port            |
| Carsten Gehrke      - Developer, Delphi (Fort Knox Project) |
| Charlie Long        - Developer, Delphi (Fort Knox Project) |
| Jon Lasser          - Original Coordinator                |
|
| and many other contributors whose names can be found at:
|
| www.bastille-linux.org/credits.html
|
| We are indebted to the following for support and help:
|
| The US TSWG and US Navy   Hewlett Packard
| Mandrakesoft              The SANS Institute
| VA Software                IBM
+-----+

```

Figure 3-29 Bastille - Author credits

31. The Bastille Linux software tends to leave artifacts on the terminal screen, so you may need to reinitialize the terminal:

```
~> clear
```

Bastille Linux has hardened your system based upon the answers to the questions it asked you. However, there is still some additional hardening to be done.

Running Bastille Linux on RHEL4

Bastille is run using an ncurses-based text interface. Run Bastille:

1. While you are still in the previously opened SSH session to the firewall, run the Bastille program:

```
~> bastille -c
```

2. Accept the Bastille license, shown in Figure 3-30 on page 131. Press the Enter key until you are given an angle bracket prompt. Type accept and press the Enter key.

Note: Remember that these settings are to be changed on the firewall partition only. Using the Bastille Linux software on the LINUXST partition could result in a system that cannot perform proper mail scanning or filtering.

```
[root@LINUXFW tools]# bastille -c
NOTE:    Using Curses user interface module.
NOTE:    Only displaying questions relevant to the current configuration.

Copyright (C) 1999-2002 Jay Beale
Copyright (C) 1999-2001 Peter Watkins
Copyright (C) 2000 Paul L. Allen
Copyright (C) 2001-2003 Hewlett-Packard Development Company, L.P.
Bastille is free software; you are welcome to redistribute it under
certain conditions.  See the 'COPYING' file in your distribution for terms.

DISCLAIMER.  Use of Bastille can help optimize system security, but does not
guarantee system security.  Information about security obtained through use of
Bastille is provided on an AS-IS basis only and is subject to change without
notice.  Customer acknowledges they are responsible for their system's security.
TO THE EXTENT ALLOWED BY LOCAL LAW, Bastille (SOFTWARE) IS PROVIDED TO YOU
AS IS WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, WHETHER ORAL OR WRITTEN,
EXPRESS OR IMPLIED.  JAY BEALE, THE BASTILLE DEVELOPERS, AND THEIR SUPPLIERS
DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.
Some countries, states and provinces do not allow exclusions of implied
warranties or conditions, so the above exclusion may not apply to you.  You may
have other rights that vary from country to country, state to state, or province
to province.  EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT WILL
JAY BEALE, THE BASTILLE DEVELOPERS, OR THEIR SUBSIDIARIES, AFFILIATES OR
SUPPLIERS BE LIABLE FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL OR OTHER
DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF
THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED
IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED
OF THE POSSIBILITY OF SUCH DAMAGES.  Your use of the Software is entirely at your
own risk.  Should the Software prove defective, you assume the entire cost of all
service, repair or correction.  Some countries, states and provinces do not allow
the exclusion or limitation of liability for incidental or consequential
damages, so the above limitation may not apply to you.

You must accept the terms of this disclaimer to use
Bastille.  Type "accept" (without quotes) within 5
minutes to accept the terms of the above disclaimer
>
```

Figure 3-30 Bastille - License

3. The Bastille user interface appears, as shown in Figure 3-31 on page 132. Tab to Next and press the Enter key.


```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to set more restrictive permissions on the administration
|utilities? [N]
|In general, the default file permissions set by most vendors are fairly secure
|. To make them more secure, though, you can remove non-root user access to
|some administrator functions.
|
|If you choose this option, you'll be changing the permissions on some common
|system administration utilities so that they're not readable or executable by
|users other than root. These utilities (which include linuxconf, fsck,
|ifconfig, runlevel and portmap) are ones that most users should never have a
|need to access. This option will increase your system security, but there's a
|chance it will inconvenience your users.
|
+-----+
|
|          +-----+
|          |Yes  |
|          |No  |
|          +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-32 Bastille - Restricting administrative functions

5. Restrict file permission, as shown in Figure 3-33.

Restricting system files to only the root user forces a would-be attacker to become the root user before changing a configuration that can be potentially damaging to your system. Becoming root is much more difficult than gaining access to a standard user account, so it is prudent to restrict file permissions. Tab to the Next button and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|The following questions all pertain to disabling "SUID root" permission for
|particular programs. This permission allows non-root users to run these
|programs, increasing convenience but decreasing security. If a security
|weakness or vulnerability is found in these programs, it can be exploited to
|gain root-level access to your computer through any user account.
|
|If you answer "Yes" and then realize later that you do need SUID permissions
|on a specific program, you can always turn it back on later with chmod u+s <
|file name>.
|
+-----+
|
|          < Back >   < Next >   < Explain Less >

```

Figure 3-33 Bastille - Restricting file permissions

6. Disable SUID for the **mount** and **umount** commands, as shown in Figure 3-34.

This forces an attacker to become root before they can mount or unmount any file systems on your server. Use the arrow keys to select Yes and press the Enter key.

Note: Forcing everyone to become the root user before they can use any kind of tool that may damage your system provides you with user accountability. Direct login to your system as the root user will be disabled later on in the Bastille program. This forces a user to log in as a local standard user before becoming root. Any damage done as root can then be traced back to the standard user account that an attacker initially used to login to the system. With knowledge of this, most users will opt to keep their passwords a secret. In the event that your system is compromised, knowing which user account was used to start with provides you with a starting point for your investigation into the attack.

```
+-----+
|                                     |
|                               Bastille |
|-----+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for mount/umount?
|Mount and umount are used for mounting (activating) and unmounting (
|deactivating) drives that were not automatically mounted at boot time. This
|can include floppy and CD-ROM drives. Disabling SUID would still allow anyone
|with the root password to mount and unmount drives.
|
|Would you like to disable SUID status for mount/umount?
|
+-----+
|                                     |
|                               +-----+ |
|                               |Yes | |
|                               |No  | |
|                               +-----+ |
|                                     |
| < Back >   < Next >   < Explain Less > |
+-----+
```

Figure 3-34 Bastille - Disabling SUID for the mount and umount commands

7. Disable SUID for the **ping** command, as shown in Figure 3-35 on page 135.

This forces an attacker to become root before they can use the **ping** command, as the **ping** command uses root privileges to open a raw socket. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for ping? [Y]
|Ping is used for testing network connectivity. Specifically it's for testing
|the ability of the network to get a packet from this machine to another and
|back. The ping program is SUID since only the root user can open a raw socket
|. Since, however, it is often used only by the person responsible for
|networking the host, who normally has root access, we recommend disabling SUID
|status for it.
|
|Would you like to disable SUID status for ping? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-35 Bastille - Disabling SUID for the ping command

8. Disable SUID for the **at** command, as shown in Figure 3-36.

This forces standard users to use the cron system to schedule jobs, so the at scheduling system cannot be exploited. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for at? [Y]
|"at" is used for scheduling an individual task to run at a single later time.
|There have historically been many exploits that take advantage of weaknesses
|in "at". Virtually all of the necessary functionality of "at" can be found in
|cron (and removing cron is not practical) so there is no need to retain
|privileged access for "at".
|
|Would you like to disable SUID status for at? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-36 Bastille - Disabling SUID for the at command

9. Disable the BSD r-tools, as shown in Figure 3-37 on page 136.


```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for usernetctl? [Y]
|usernetctl is a utility that allows ordinary users to control the network
|interfaces. In general, there's no reason for anyone other than the system
|administrator to control network interfaces.
|
|Would you like to disable SUID status for usernetctl? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-38 Bastille - Disabling SUID for usernetctl

11. Disable SUID for the traceroute program, as shown in Figure 3-39.

Traceroute is useful for debugging network issues. This is not necessary for non-root users. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+FilePermissions.pm Module 2 of 0-----+
|Q: Would you like to disable SUID status for traceroute? [Y]
|The traceroute utility is used to test network connectivity. It is useful for
|debugging network problems, but it is generally not necessary, especially for
|nn-privileged users. If non-root users will be needing to debug network
|connections, you can leave the SUID bit on traceroute. Otherwise, you should
|disable it.
|
|Would you like to disable SUID status for traceroute? [Y]
|
+-----+
|
| +-----+
| |Yes |
| |No |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-39 Bastille - Disabling SUID for traceroute

12. Disable clear-text r-protocols, as shown in Figure 3-40 on page 138.


```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Would you like to enforce password aging? [Y]
|Your operating system's default behavior, which we would change here, is to
|disable an account when the password hasn't changed in 99,999 days. This
|interval is too long to be useful. We can set the default to 60 days. At
|some point before the 60 days have passed, the system will ask the user to
|change his or her password. At the end of the 60 days, if the password has
|not been changed, the account will be temporarily disabled. We'll make sure
|this warning period is at least 5 days long. We would make this change in /
|etc/login.defs.
|
|Would you like to enforce password aging? [Y]
|
+-----+
|
| +-----+
| | Yes |
| | No  |
| |     |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-41 Bastille - Leaving the password aging settings

14. Set a default umask, as shown in Figure 3-42.

The umask is used to set file permissions that users create. We recommend that users only be allowed to read and write to their own files. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Do you want to set the default umask? [Y]
|The umask sets the default permission for files that you create. Bastille can
|set one of several umasks in the default login configuration files. These
|cover standard shells like csh and most bourne shell variants like bash, sh,
|and ksh. If you are going to install other shells, you may have to configure
|them yourself. The only reason not to set at least a minimal default umask is
|if you are sure that you have already set one.
|
|Do you want to set the default umask? [Y]
|
+-----+
|
| +-----+
| | Yes |
| | No  |
| |     |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-42 Bastille - Setting a default umask


```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Should we disallow root login on tty's 1-6? [N]
|You can restrict which tty's root can login on. Some sites choose to restrict
|root logins, so that an admin must login with an ordinary user account and
|then use su to become root.
|
|This can stop an attacker who has only been able to steal the root password
|from logging in directly. He has to steal a second account's password to make
|use of the root password via the ttys.
|
|Should we disallow root login on tty's 1-6? [N]
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-44 Bastille - Restricting root login on system consoles

17. Remove extraneous accounts, as shown in Figure 3-45.

These accounts are shipped with the operating system. However, they are not necessary for the proper functionality of the firewall. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Should Bastille ask you for extraneous accounts to delete?
|Most operating systems ship with a number of accounts that are extraneous or
|at least not used by systems that have a specific purpose. Bastille can remove
|extraneous accounts from the system. If you choose Yes, the next question
|will ask you for a list of accounts and will recommend a list to you.
|
|Should Bastille ask you for extraneous accounts to delete?
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-45 Bastille - Removing extraneous accounts

18. The extraneous accounts are listed, as shown in Figure 3-46. The default accounts listed are acceptable choices to remove. Tab to the Next button and press the Enter key.

```
+-----+
|                                     |
|                                     | Bastille                               |
|                                     |
+-----+
+AccountSecurity.pm Module 3 of 0-----+
|Q: Which extraneous accounts should Bastille delete (space-separated) ?
|Most operating systems ship with a number of accounts that are extraneous or
|at least not used by systems that have a specific purpose. Bastille can remove
|extraneous accounts from the system.
|
|Please specify a list of accounts to delete in a space-separated list.
|
|Samples follow:
|Red Hat Enterprise Linux 3: gopher, games SuSE Enterprise 9: games, uucp
|Which extraneous accounts should Bastille delete (space-separated) ?
+-----+
|Answer: games gopher
+-----+
|                                     |
|                                     | < Back > < Next > < Explain Less > |
|                                     |
```

Figure 3-46 Bastille - Extraneous accounts to remove

19. Disable Ctrl+Alt+Delete rebooting, as shown in Figure 3-47 on page 143.

Note: The Bastille program states that unless the powerline, power switch, and case of the server can be physically protected, disabling the ability to reboot via the Ctrl+Alt+Delete sequence is unnecessary. However, it is most likely that an attacker will not have physical access to the server. Disabling the Ctrl+Alt+Delete sequence reduces the risk of an attacker being able to reset the server.

There is a good chance that an attacker would be using a virtual console to attack your system, in which case disabling the Ctrl+Alt+Delete reboot is a prudent security measure. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Would you like to disable CTRL-ALT-DELETE rebooting? [N]
|Disabling CTRL-ALT-DELETE rebooting is designed to prevent an attacker with
|access to the machine's keyboard from being able to reboot the machine. A
|reboot done in this manner should not damage the file system, as it shuts the
|machine down cleanly, writing out all pending data in the disk cache to disk
|first. Even with this functionality disabled, however, an attacker could just
|power cycle machine or pull the power cord.
|
|Unless the power line, switch and case of the machine can be physically
|protected, this precaution is wholly unnecessary. Given the fact that the
|attacker _can_ reboot the machine, would you prefer that s/he do it in a way
|potentially damages the file system? Think carefully here, as maintaining the
|integrity of the machine's file system may be secondary to the goal of keeping|
+-----+
|
|          +-----+
|          |Yes  |
|          |No  |
|          +-----+
|
| < Back >  < Next >  < Explain Less >

```

Figure 3-47 Bastille - Disabling Ctrl+Alt+Delete rebooting

20. Password protect single user mode, as shown in Figure 3-48.

Single user mode allows full root privileges to anyone who has access to an unprotected boot loader. Password protecting the boot loader stops this attack. Use the arrow keys to select Yes and press the Enter key.

```

+-----+
| Bastille |
+-----+
+BootSecurity.pm Module 4 of 0-----+
|Q: Would you like to password protect single-user mode? [Y]
|Anyone who can physically interact with your system can tell the bootloader to
|bring your machine up in "single user mode", where s/he is given root
|privileges and everyone else is locked out of the system. This doesn't
|require a password on most Unix systems. The method differs with the
|bootloader being used, thus on each operating system revision and
|architecture. You can test this attack on a Linux system that uses LILO by
|typing "linux single" at the LILO: prompt.
|
|Bastille can password-protect the bootprompt for you. You won't have to
|remember another password--single user mode, or "root" mode, will require the
|root password.
|
+-----+
|
|          +-----+
|          |Yes  |
|          |No  |
|          +-----+
|
| < Back >  < Next >  < Explain Less >

```

Figure 3-48 Bastille - Password protecting single user


```

+-----+
| Bastille |
+-----+
+SecureInetd.pm Module 5 of 0-----+
|Q: Would you like to display "Authorized Use" messages at log-in time? [Y]
|At this point you can create "Authorized Use Only" messages for your site.
|These may be very helpful in prosecuting system crackers you may catch trying
|to break into your system. Bastille can make default messages which you may
|then later edit. This is sort of like an "anti-welcome mat" for your computer
|.
|
|Would you like to display "Authorized Use" messages at log-in time? [Y]
|
+-----+
| Yes |
| No  |
+-----+
< Back >  < Next >  < Explain Less >

```

Figure 3-50 Bastille - Creating an authorized use message

23. Customize the authorized use message by typing the name or e-mail address of the system administrator, as shown in Figure 3-51 on page 146. You may instead type the name of your organization.

We recommend that you include some method of communication to the system administrator, as it cuts off the argument that an attacker was unable to ask for permission to use the system, because there was no way to contact the system administrator. Use Tab to select the Next button and press the Enter key.

```
+-----+
| Bastille |
+-----+
+SecureInetd.pm Module 5 of 0-----+
|Q: Who is responsible for granting authorization to use this machine?
|Bastille will start to make the banner more specific by telling the user who
|is responsible for this machine. This will state explicitly from whom the
|user needs to obtain authorization to use this machine. Please type in the
|name of the company, person, or other organization who owns or is responsible
|for this machine.
|
|Who is responsible for granting authorization to use this machine?
|
+-----+
| Answer: Alex Robar |
+-----+
| < Back > < Next > < Explain Less > |
+-----+
```

Figure 3-51 Bastille - Customizing the authorized use message

24. Limit system resources, as shown in Figure 3-52 on page 147.

This partition will only be used to protect your network and detect network intrusions. As such, it will never require any user to be running more than 150 simultaneous processes. Use the arrow keys to select Yes and press the Enter key.

You will be informed that the system resource limits have been set in the file /etc/security/limits.conf. Press the Tab key to continue.

```

+-----+
| Bastille |
+-----+
+ConfigureMiscPAM.pm Module 7 of 0-----+
|Q: Would you like to put limits on system resource usage? [N]
|Denial of Service attacks are often very difficult to defend against, since
|they don't require access of any kind to the target machine. Since several
|major daemons, including the web, name, and FTP servers, may run as a
|particular user, you can limit the effectiveness of many Denial of Service
|attacks by modifying /etc/security/limits.conf. If you restrict the resources
|available in this manner, you can effectively cripple most Denial of Service
|attacks.
|
|If you choose this option, you'll be setting the following initial limits on
|resource usage:
|
| - The number of allowed core files will be set to zero. Core files
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-52 Bastille - Limiting system resources

25. Do not restrict console login privileges to a small group of users, as shown in Figure 3-53.

It is feasible for this to be defined; however, the setup for the firewall partition should only include one standard user in addition to the root user. Adding additional users is a security risk, and as such, there should be no need to define which users can log in at a console. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+ConfigureMiscPAM.pm Module 7 of 0-----+
|Q: Should we restrict console access to a small group of user accounts? [N]
|Under some distributions, users logged in at the console have some special
|access rights (like the ability to mount the CD-ROM drive). You can disable
|this special access entirely, but a more flexible option is to restrict
|console access to a small group of trusted user accounts.
|
|Should we restrict console access to a small group of user accounts? [N]
|
+-----+
|
| +-----+
| |Yes |
| |No  |
| +-----+
|
| < Back >   < Next >   < Explain Less >

```

Figure 3-53 Bastille - Leaving console group login privileges


```
+-----+
| Bastille |
+-----+
+MiscellaneousDaemons.pm Module 9 of 0-----+
|To make the operating system more secure, we try to deactivate all system
|daemons, especially those running at a high/unlimited level of privilege.
|Each active system daemon serves as a potential point of break-in, which might
|allow an attacker illegitimate access to your system. An attacker can use
|these system daemons to gain access if they are later found to have a bug or
|security vulnerability.
|
|We practice a minimalist principle here: minimize the number of privileged
|system daemons and you can decrease your chances of being a victim should one
|of the standard daemons be found later to have a vulnerability. This section
|will require careful attention, but if you have doubts, you should be able to
|safely select the default value in most cases.
|
+-----+

< Back > < Next > < Explain Less >
```

Figure 3-55 Bastille - Securing system daemons

28. Disable the Advanced Power Management™ Daemon (APMD), as shown in Figure 3-56.

APMD is not used by any Linux operating system that is installed into an LPAR. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
| Bastille |
+-----+
+MiscellaneousDaemons.pm Module 9 of 0-----+
|Q: Would you like to disable apmd? [Y]
|apmd is used to monitor battery power and is used almost exclusively by
|notebook/laptop computers.
|
|Would you like to disable apmd? [Y]
|
+-----+

< Back > < Next > < Explain Less >
```

Figure 3-56 Bastille - Disabling APMD

29. Disable the Personal Computer Memory Card International Association (PCMCIA) services, as shown in Figure 3-57 on page 150.

PCMCIA interfaces are primarily used in laptops. There are not used in the eServer i5. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
| Bastille |
+-----+
+MiscellaneousDaemons.pm Module 9 of 0-----+
|Q: Would you like to disable PCMCIA services? [Y]
|If this machine is not a notebook, it probably has no PCMCIA ports.  PCMCIA
|ports allow the use of easily removable credit-card-sized devices.  If this
|machine has no PCMCIA ports, you should be able to disable PCMCIA services
|without any problems.
|
|Would you like to disable PCMCIA services? [Y]
|
+-----+
| < Back >   < Next >   < Explain Less > |
+-----+
```

Figure 3-57 Bastille - Disabling PCMCIA services

30. Disable the General Purpose Mouse (GPM) server, as shown in Figure 3-58.

There is no reason for a mouse to be used for system administration. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
| Bastille |
+-----+
+MiscellaneousDaemons.pm Module 9 of 0-----+
|Q: Would you like to disable GPM? [Y]
|GPM is used in console (text) mode to add mouse support to text mode.  If you
|will be using this machine in console mode and will want mouse support, leave
|GPM on.
|
|Would you like to disable GPM? [Y]
|
+-----+
| < Back >   < Next >   < Explain Less > |
+-----+
```

Figure 3-58 Bastille - Disabling the GPM server

31. Disable the hardware detection service, Kudzu, at boot, as shown in Figure 3-59 on page 151.

New hardware is rarely added to servers. If hardware is added, a system administrator can easily run the Kudzu service manually. If an attacker happens to have physical access

to the system, it is possible to introduce corrupt or virus ridden files into the system via external storage devices that are loaded at boot if the Kudzu service runs automatically. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
|                                     |
|                                     | Bastille                               |
|                                     |-----+
|+MiscellaneousDaemons.pm Module 9 of 0-----+
|{Q: Would you like to deactivate kudzu's run at boot?
|{The kudzu hardware detection daemon, created by Red Hat, runs on each boot,
|{checks for new hardware, helps configure it if present, and then terminates.
|{This can be a very useful daemon on workstation machines where users change
|{their own hardware frequently. On the other hand, this daemon can allow
|{unprivileged users (non-system administrators) to add and configure hardware
|{with full root privilege. This generates some additional risk.
|
|{We believe that few environments need to give ordinary users this kind of
|{privilege. This program can be safely deactivated. Even after such
|{deactivation, sysadmins can indeed run kudzu from the command line to get the
|{very same hardware detection and configuration functionality.
|
|-----+
|
|                                     |
|                                     | < Back >   < Next >   < Explain Less > |
|                                     |-----+
+-----+
```

Figure 3-59 Bastille - Disabling Kudzu at boot

32. Disable the sendmail daemon, as shown in Figure 3-60.

Any incoming mail will simply be forwarded to the LINUXST partition. Sendmail is not used on the LINUXFW partition. Use the arrow keys to select Yes and press the Enter key.

```
+-----+
|                                     |
|                                     | Bastille                               |
|                                     |-----+
|+MiscellaneousDaemons.pm Module 9 of 0-----+
|{Q: Do you want to stop sendmail from running in daemon mode? [Y]
|{You do not need to have sendmail running in daemon mode to send and receive
|{email, and unless you have a constant network connection, you probably cannot
|{run sendmail in daemon mode. Daemon mode means that sendmail is constantly
|{listening on a network connection waiting to receive mail.
|
|{If you disable daemon mode, Bastille will ask you if you would like to run
|{sendmail every few minutes to process the queue of outgoing mail. Most
|{programs which send mail will still do so immediately, and processing the
|{queue will take care of transient errors.
|
|{If you receive all of your email via a POP/IMAP mailbox provided by your ISP
|{you may have no need of daemon-mode sendmail, unless you're running a
|
|-----+
|
|                                     |
|                                     | < Back >   < Next >   < Explain Less > |
|                                     |-----+
+-----+
```

Figure 3-60 Bastille - Disabling the Sendmail daemon

33. Bastille shows you a recommendation for changes to the Apache Web server configuration, as seen in Figure 3-61.

As the firewall partition will not be running any Web server, this can be safely ignored. Tab to the Next button and press the Enter key to continue.

```
+-----+
|                               Bastille                               |
+-----+
+Apache.pm Module 12 of 0-----+
| There are a few other changes that we recommend you make to the web server's |
| configuration. There are very few intrinsic security flaws in the Apache web  |
| server, but there are two important ones:                                     |
|                                                                              |
|   As with all web servers, it is generally required to send and receive    |
|   information to and from anyone on the internet.                          |
|                                                                              |
|   In many environments, the people telling the server how to behave are not  |
|   knowledgeable system administrators by trade. Before you discount this fact |
|   , take account of the wide proliferation of configurations under which any |
|   user on the system can instruct the server to execute arbitrary code for   |
|   anyone who comes to the site, via CGI scripts.                          |
+-----+

< Back >   < Next >   < Explain Less >
```

Figure 3-61 Bastille - Leaving Apache server settings

34. Do not disable printing, as seen in Figure 3-62 on page 153.

Printing will not be used by default. However, it may be useful for having a hard copy of logs should they be needed. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+Printing.pm Module 13 of 0-----+
|Q: Would you like to disable printing? [N]
|If this machine is not going to need to print, you should stop the print
|scheduler and restrict the permissions on all the printing utilities.
|
|On Linux, you could revert this later by typing:
|
|# /bin/chmod 0755 /usr/bin/lpr /usr/bin/lprm /usr/bin/lpstat # /bin/chmod
|04755 /usr/bin/lppasswd
|
|# /sbin/chkconfig cups on
|
|This is only recommended if this machine will not be used for printing in the
|near future. If you deactivate this, you might want to write down the
+-----+

< Back > < Next > < Explain Less >

```

Figure 3-62 Bastille - Leaving printing daemon settings

35. Do not enable TMPDIR scripts, as shown in Figure 3-63.

The security of the /tmp file system will be addressed later in 3.1.3, “Altering insecure defaults” on page 156. Enabling the TMPDIR scripts could interfere with this security measure. Use the arrow keys to select No and press the Enter key.

```

+-----+
| Bastille |
+-----+
+TMPDIR.pm Module 17 of 0-----+
|Q: Would you like to install TMPDIR/TMP scripts? [N]
|Many programs use the /tmp directory in ways that are dangerous on multi-user
|systems. Many of those programs will use an alternate directory if one is
|specified with the TMPDIR or TMP environment variables. We can install scripts
|that will be run when users log in that safely create suitable temporary
|directories and set the TMPDIR and TMP environment variables. This depends on
|your system supporting /etc/profile.d scripts.
|
|Would you like to install TMPDIR/TMP scripts? [N]
+-----+

+-----+
| Yes |
| No  |
+-----+

< Back > < Next > < Explain Less >

```

Figure 3-63 Bastille - Leaving TMPDIR script settings


```

+-----+
| Bastille |
+-----+
+End of 0-----+
|Q: Are you finished answering the questions, i.e. may we make the changes?
|We will now implement the choices you have made here.

|Answer NO if you want to go back and make changes!

|Are you finished answering the questions, i.e. may we make the changes?

+-----+
| Yes |
| No  |
+-----+
< Back >   < Next >   < Explain Less >

```

Figure 3-65 Bastille - Ready to make changes

38. You will be presented with the credits of the Bastille Linux team, as seen in Figure 3-66. Press the Tab key to continue.

```

+Bastille Credits (press TAB to go on)-----+
| Jay Beale - Lead Architect and Original Author |
| HP Bastille Dev Team - Developers - HP-UX Port, Design/Arch. |
| Peter Watkins - Core Developer: Firewall |
| Mike Rash - Developer: PSAD |
| Paul Allen - Developer: User Interface |
| Javier Fernandez-Sanguino - Developer - Debian Port |
| Niki Rahimi (IBM) - Developer - SuSE and TurboLinux Ports |
| Brian Stine - Developer - Gentoo Port |
| Carsten Gehrke - Developer, Delphi (Fort Knox Project) |
| Charlie Long - Developer, Delphi (Fort Knox Project) |
| Jon Lasser - Original Coordinator |
|
| and many other contributors whose names can be found at: |
|
| www.bastille-linux.org/credits.html |
|
| We are indebted to the following for support and help: |
|
| The US TSWG and US Navy Hewlett Packard |
| Mandrakesoft The SANS Institute |
| VA Software IBM |
+-----+

```

Figure 3-66 Bastille - Author credits

39. The Bastille Linux software tends to leave artifacts on the terminal screen, so you may need to reinitialize the terminal:

```
~> clear
```

Bastille Linux has hardened your system based upon the answers to the questions it asked you. However, there is still some additional hardening to be done.

3.1.2 Removing unnecessary servers

SLES9 and RHEL4 both come with a few servers that are not necessary for the firewall partition to function. We recommend that anything that is not necessary for functionality be disabled or removed.

To remove unnecessary servers in SLES9:

```
~> rpm -e yast2-dns-server-2.9.24-0.2
~> rpm -e yast2-dhcp-server-2.9.23-0.2
~> rpm -e yast2-http-server-2.9.26-1.2
~> rpm -e yast2-ntp-client-2.9.14-0.3
~> rpm -e yast2-tftp-server-2.9.4-23.1
```

To remove unnecessary servers in RHEL4:

```
~> rpm -e ftp-0.17-22
~> rpm -e lftp-3.0.6-3
```

Unnecessary servers have been removed from your system.

3.1.3 Altering insecure defaults

While most installations of SLES9 and RHEL4 are moderately secure, the default settings are for a system that is intended to be a multi-user server setup. The firewall partition does not need to have multiple users log on to it, and it will not have a very heavy processing load. As such, there are a few default settings that can be changed to add additional security. The following steps must be completed:

- ▶ Move the /tmp file system off the hard drive.
- ▶ Display the authorized use message at all login prompts.
- ▶ Restrict system account logins.
- ▶ Disable TCP SYN and IP spoofing attacks.
- ▶ Disable SSH1.

This section contains five subsections that provide instructions for completing each one of these tasks.

Move /tmp file system

The /tmp file system contains many dozens of files that are either being written to disk or accessed from disk. By default, it is stored on the physical hard disk. As a security precaution, the /tmp file system should be moved to the RAM, as this will cause the files to be overwritten more often. In addition, in the event of a reboot, the files are truly temporary, as the RAM is cleared upon a powerdown of the system. Moving the /tmp file system to the RAM reduces the threat of anyone being able to read sensitive information that has not been cleared from the system yet. To move the /tmp file system onto the RAM, follow these instructions:

1. Connect to the virtual console of your firewall system.

Note: The virtual console must be used, as moving the /tmp file system requires the system to be in single-user mode. A user connected to a system via SSH will be disconnected when that system drops out of multi-user mode.

- a. Open up the PuTTY client.
 - b. Enter your i5/OS server name as the Host Name (or IP address).
 - c. Set the protocol to Telnet.
 - d. Enter 2301 as the Port.
 - e. Click the **Open** button to open the connection.
2. At the Linux shell, drop to single-user mode:

```
~> init 1
```
 3. When you are prompted, enter the root password from cell J2 on the planning worksheet, and press the Enter key.
 4. Clear the contents of the existing /tmp file system:

```
~> rm -r /tmp/.[a-zA-Z]*
```
 5. Open /etc/fstab for editing in the vi text editor:

```
~> vi /etc/fstab
```
 6. The file system table file will open for editing using the vi text editor. Press the Insert key to enter editing mode, and enter the following line after the last line in the file:

```
none /tmp tmpfs noexec,nosuid,nodev 1 2
```

This line tells Linux to mount the /tmp file system onto the RAM, as opposed to a location on a local disk.
 7. Press the Esc button to exit editing mode.
 8. Type :wq to save the file and exit the editor.
 9. At the Linux command shell, mount the newly created /tmp file system:

```
~> mount -a
```
 10. Return to multi-user mode:

```
~> init 3
```
 11. Close the PuTTY window, as the remainder of the system administration will be done via SSH.

Displaying an authorized use message

The authorized use message displays a warning about who your system belongs to, and who is authorized to use the system. The message is displayed upon login to any console. However, the message is not displayed by default when a user connects to your system via SSH. To display the authorized use message to all users who connect to your system via SSH, follow these instructions:

1. Open up an SSH connection to the firewall partition:
 - a. Open up the PuTTY client.
 - b. Enter the IP address of the LINUXFW partition from cell CB2 on the planning worksheet in the Host Name (or IP address) field.
 - c. Set the protocol to SSH.
 - d. Click the **Open** button to open the connection.

2. Log in as the root user, using the password from cell J2 of the planning worksheet.
3. Copy the authorized use message created by Bastille to the Message of the Day (MOTD) file so that SSH users will see it, too:

```
~> cp /etc/issue /etc/motd
```

If you are prompted to confirm overwriting an existing file, type *y* and press the Enter key.

Restrict system logins

A system account is an account installed by default on a Linux operating system that has a specific role, such as system logging. The account may have privileges that are higher than that of a normal user. Follow these steps to only allow login to your system as a regular user:

1. In the SSH shell that you opened earlier, open the SSH daemon configuration file for editing in the vi text editor:
2. You will be presented with a vi text editor that has the `sshd` file open for editing. Press the Insert key, and enter the following line below the last line in the file:

```
account    required    /lib/security/pam_access.so
```

3. Press the Esc key to exit editing mode.
4. Type `:wq` to save the file and quit the text editor.
5. Open the `access.conf` file for editing in the vi text editor:
6. You will be presented with a vi text editor that has the `access.conf` file open for editing. Page down to the bottom of the file. Press the Insert key, and enter the following line:

```
 -:ALL EXCEPT users :ALL
```

This line tells the system that no users are to be granted logon to the system unless they are contained within the *users* group.

7. Press the Esc key to exit editing mode.
8. Type `:wq` to save the file and quit the text editor.

Tune TCP parameters

A SYN flood is an attack that uses the SYN TCP flagset to send requests to the server faster than it can process them. Each SYN flood packet appears to come from a different IP address, all of which have been spoofed. A SYN flood attack can cause a server to crash rapidly. Protect against SYN flood and IP spoofing attacks by following these steps:

1. Open `sysctl.conf` for editing in the vi text editor:
2. You will be presented with a vi text editor that has a blank file open for editing. Press the Insert key, and enter the following line:

```
net.ipv4.tcp_syncookies = 1
```

This line protects against SYN flood attacks.

3. Press the Enter key to move the cursor to a new blank line, and enter the following:

```
net.ipv4.conf.all.rp_filter = 1
```

This line protects against IP spoofing.

4. Exit editing mode by pressing the Esc key.
5. Save the file by typing `:wq`.

6. Activate the new additions to the `sysctl.conf` file:

```
~> sysctl -p
```

Disable SSH1

SSH is a secure protocol for connecting to and using remote server systems. Like most software packages, SSH has been updated over time to include improvements in functionality and security. The newer SSH protocol, SSH2, is not vulnerable to any of the security issues that may have compromised a system using SSH1. Configure SSH to increase security by disabling the older, less secure protocol, SSH1:

1. In the SSH shell to the firewall partition that you opened earlier, open the SSH daemon configuration file for editing in the vi text editor:

```
~> vi /etc/ssh/sshd_config
```

2. Press the Insert button to enter into editing mode.
3. Alter the line that reads:

```
#Protocol 2,1  
to read:  
Protocol 2
```

4. Press the Esc button to exit editing mode.
5. Type `:wq` to save the file and exit the vi text editor.
6. Open the SSH client configuration file for editing in the vi text editor:

```
~> vi /etc/ssh/ssh_config
```

Note: Ensuring that outgoing SSH from your system can only use the SSH2 protocol is a prudent security measure, as it simply ensures that any data leaving your system for a remote destination is as secure as it can be between those two points.

7. Press the Insert button to enter into editing mode.
 8. Alter the line that reads:
- ```
Protocol 2,1
to read:
Protocol 2
```
9. Press the Esc button to exit editing mode.
  10. Type `:wq` to save the file and exit the vi text editor.
  11. Restart the SSH daemon:

```
~> /etc/init.d/sshd restart
```

Insecure defaults have now been corrected. The firewall partition now only allows the more secure SSH2 protocol to be used; the `/tmp` file system is loaded onto RAM, making it truly temporary; and direct login as root is entirely disabled from all login points.

## 3.2 iptables rules

This section details setting up the firewall partition to block all traffic that could be potentially damaging to your system. Notice in Figure 2-1 on page 21 that even traffic from your local network will go through the firewall and Snort intrusion detection systems. This ensures that your servers are protected not only from external sources, such as malicious hackers, but also

from seemingly benign internal sources. With increased mobility and remote links, it is much easier for someone to unknowingly bring the latest worm or Trojan horse into your network. Forcing all traffic through the security systems ensures that your servers will remain intact even if a trusted system introduces a threat to your network.

Configuration of the firewall is one of the most important steps in making the system secure. The Linux operating system comes with built-in packet filtering features called netfilter/iptables. Netfilter and iptables work in conjunction with one another to form the building blocks of a rock solid firewall system. The iptables system is the set of rules that must be defined to determine how to handle any incoming or outgoing network packets.

It is likely that the two Linux servers set up in this redpaper will not be the sole servers on your network. In this case, it is important that you are able to construct your own set of iptables rules to add to the default set provided for you in this redpaper. As such, in this section we discuss how to set up iptables rules, in addition to providing a good set of base rules for any system.

### 3.2.1 Understanding iptables

Before an example of a firewall is given, it is a good idea to understand how the iptables firewall system works.

Iptables breaks rulesets down into three levels:

#### ► Tables

Tables are rulesets with a specific desired outcome. There are three built-in tables in the iptables system:

- Filter: The filter table accepts or denies packets based upon a definable ruleset.
- Network Address Translation (NAT): The NAT table translates the source or destination address of a packet based upon a definable ruleset.
- Mangle: The mangle table alters other aspects of a packet based upon a definable ruleset.

#### ► Chains

Chains are defined rule paths that packets follow. The filter table has the INPUT, FORWARD, and OUTPUT chains built in. The INPUT chain receives packets that your system is receiving. The FORWARD chain receives packets that are simply being routed through your system. The OUTPUT chain receives packets that are leaving your system.

#### ► Rules

Rules are the heart of iptables. They interpret a packet based upon header and content information. If a packet matches a rule, an appropriate action is performed based upon the defined action of the parent chain. Actions include dropping, accepting, or altering packets.

This redpaper includes information about setting up the filter table using both the built-in and custom chains.

### Chains

Chains in iptables are quite simply multiple rules that are chained together. Chains may be used to apply the same action to different types of packets.

This section provides some of the most commonly used chain options. To create a new chain with the name *chainname*, use the following command:

```
~> iptables -N chainname
```

To delete a chain with the name *chainname*, use the following command:

```
~> iptables -X chainname
```

To change the action initiated by a chain when a match is found to a contained rule, use the `-P` switch. To change a chain with the name *chainname* to accept a packet when a rule match is found, use the following command:

```
~> iptables -P chainname ACCEPT
```

To change the same chain to drop any packet that matches a rule, use the following command:

```
~> iptables -P chainname DROP
```

To print out the current list of rules contained within the *chainname* chain, enter the following command:

```
~> iptables -L chainname
```

To flush all the rules out of *chainname*, use the following command:

```
~> iptables -F chainname
```

To reset the packet and byte counter for *chainname*, use the following command:

```
~> iptables -Z chainname
```

**Attention:** If no chain name is specified when using the `-X` or `-F` switches, the command will be executed on *all* chains. If you do not specify a chain name, all rules will be lost, or all custom chains will be deleted.

## Rules

Rules are patterns that iptables uses to see if action needs to be taken on a specific packet. Rule syntax is much more complicated than that of chain syntax, as there are many more parameters per rule.

To append a rule to the filter table into the built-in INPUT chain, use the following command:

```
~> iptables -A INPUT rule
```

To delete a rule from the INPUT chain in the filter table, use the following command:

```
~> iptables -D INPUT rule
```

To replace a rule in the INPUT chain in the filter table, use the following command:

```
~> iptables -R INPUT rule_number rule
```

To append a rule in the INPUT chain in the NAT table, use the following command:

```
~> iptables -t nat -A INPUT rule
```

The following generic switches allow you to define your rule:

- ▶ `-p`: Protocol. Used to check if a packet is using a specified protocol. Common protocols are TCP, UDP, and ICMP. Both SLES9 and RHEL4 support upwards of 130 protocols by default, a complete list of which can be found in the `/etc/protocols` file.

- ▶ **-s:** Source. Used to match the source of the inspected packet. The source IP address should be specified. Any packets transmitted from the given IP will count as a match.
- ▶ **-d:** Destination. Used to match the destination of the inspected packet. The destination IP address should be specified. Any packets with a final destination of the given IP will count as a match.
- ▶ **-i:** Incoming interface. The network interface from which the packet was received. Note that this switch can only be used with the FORWARD and INPUT chains. Interfaces are specified by their system identifier, such as eth0.
- ▶ **-o:** Outgoing interface. The network interface that the packet is using to leave the network. Note that this switch can only be used with the FORWARD and OUTPUT chains. Interfaces are specified by their system identifier, such as eth0.
- ▶ **-m state:** Match state. The state that any given packet has, as defined by the iptables connection state tracker. The -m state option will be followed by the --state *state* flag, where *state* can be ESTABLISHED, RELATED, NEW, or INVALID.

The following switches are specific to the TCP and UDP protocols. Your rule must have the -p TCP or -p UDP switches for these switches to be recognized:

- ▶ **--sport:** Source port. The port that the packet originated from. Ports are specified numerically, such as 22. A range of ports can be specified using a colon. To specify ports 22 through 80, you used 22:80.
- ▶ **--dport:** Destination port. The port that the packet is destined to. Ports are specified numerically. A range of ports may also be specified in the same fashion as the --sport switch.
- ▶ **--tcp-flags / --udp-flags:** Flags specific to TCP or UDP packets, such as SYN (signal new connection), ACK (acknowledge data receipt), FIN (close the current connection), or RST (drop the current connection).

**Note:** All of these rules can be negated using an exclamation point (!). For example, if you wish to match all packets that are *not* coming from a specific IP address, you would simply use the following command:

```
~> iptables -A INPUT -s ! ip_address
```

The last commonly used switch is the jump, or -j switch. The jump switch allows you to perform an action before going to the next rule. The parameter for jump may be ACCEPT, REJECT, or DROP, in which case the packet would follow the given action and stop traversing whatever chain it is in. The parameter may also be a target chain. If this is the case, the packet will then traverse the target chain. If it reaches the end of the target chain with no matches, it will continue to traverse the parent chain it came from.

**Note:** This is a small subset of the available chain and rule options. These options should give you a good start to creating a basic security policy. Advanced routing, NATing, and filtering scenarios may require additional switches.

### 3.2.2 Initial iptables setup

The basic policy that we implement provides good security measures against both internal and external threats. All incoming mail messages are forwarded to the LINUXST partition.

The rules will be placed into an executable script to be run at startup. Follow these steps to create a basic security policy:

1. Create a file that will load the iptables rules at boot time:

```
~> vi /etc/firewall.start
```

2. Press the Insert key to enter editing mode.

3. Enter the following lines:

```
#!/bin/sh
firewall.start
version=1.0
name="firewall.start"

echo "\nStarting firewall and netfilter logging"
echo "\nUsing $name script v.$version."
echo "\nSecuring Your E-mail Server on Linux redpaper\n"

This section enables IP forwarding, as LINUXST and the
i5/OS connect to the LAN and internet through the firewall.
echo " Enabling forwarding : "
echo "1" > /proc/sys/net/ipv4/ip_forward
echo " Done.\n"

This section removes all existing rules and custom chains.
echo " Flushing any existing rules..."
iptables -F INPUT
iptables -F OUTPUT
iptables -F FORWARD
iptables -t nat -F
echo " Done.\n"

This section sets the default policy on the INPUT, OUTPUT
and FORWARD chains. The policies allow for all packets to
exit the firewall, but only those which are explicitly allowed
to enter in.
echo " Setting default policy on INPUT, OUTPUT and FORWARD..."
iptables -P INPUT DROP
iptables -P OUTPUT ACCEPT
iptables -P FORWARD ACCEPT
echo " Done.\n"

This section allows all traffic on the lo interface.
echo " Allowing everything on the loopback interface..."
iptables -A INPUT -i lo -j ACCEPT
iptables -A OUTPUT -o lo -j ACCEPT
echo " Done.\n"

This section blocks XMAS and NULL packets. XMAS and NULL packets
are often generated by popular vulnerability scanning tools, and
have no valid uses, so it is safe to drop them.
echo " Blocking XMAS and NULL packets..."
iptables -A INPUT -p tcp --tcp-flags ALL ALL -j DROP
iptables -A FORWARD -p tcp --tcp-flags ALL ALL -j DROP
iptables -A INPUT -p tcp --tcp-flags ALL NONE -j DROP
iptables -A FORWARD -p tcp --tcp-flags ALL NONE -j DROP
echo " Done.\n"

This section sets up safe forwarding rules. All connections
that are established already, or related to an already established
connection will be forwarded.
```

```

echo " Setting safe forwarding rules..."
iptables -A FORWARD -i eth0 -o eth2 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -i eth0 -o eth3 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -i eth1 -o eth2 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -i eth1 -o eth3 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -i eth0 -o eth1 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -o eth0 -j ACCEPT
iptables -A FORWARD -j LOG
echo " Done.\n"

This section sets rules to accept required ports: 22 for SSH,
25 for SMTP, 80 for i5/OS web client access, 110 for POP3
mail access, 443 for secure i5/OS web client access, 763 for
SpamAssassin updates, and 6277 for e-mail blacklisting checks.
echo " Setting required ports for administration, ICMP, SMTP, mail filtering, and
POP3..."
iptables -A INPUT -p TCP --dport 22 -j ACCEPT
iptables -A INPUT -p TCP --dport 25 -j ACCEPT
iptables -A INPUT -p TCP --dport 80 -j ACCEPT
iptables -A INPUT -p TCP --dport 110 -j ACCEPT
iptables -A INPUT -p TCP --dport 443 -j ACCEPT
iptables -A INPUT -p TCP --dport 763 -j ACCEPT
iptables -A INPUT -p UDP --dport 6277 -j ACCEPT
iptables -A INPUT -p icmp -j ACCEPT
iptables -A INPUT -j LOG
echo " Done.\n"

This section sets up natting to forward incoming mail messages,
mail retrieval requests, SpamAssassin updates, SpamAssassin
blacklisting checks and i5/OS web access to the proper
partitions over virtual ethernet.
echo " Setting NAT for SMTP, POP3 and SpamAssassin updates..."
iptables -t nat -A PREROUTING -i eth0 -p TCP --dport 25 \
-j DNAT --to-destination 10.1.1.20
iptables -t nat -A PREROUTING -i eth1 -p TCP --dport 25 \
-j DNAT --to-destination 10.1.1.20
iptables -t nat -A PREROUTING -p TCP --dport 80 -j DNAT --to-destination 172.27.72.20
iptables -t nat -A PREROUTING -p TCP --dport 443 -j DNAT --to-destination 172.27.72.20
iptables -t nat -A PREROUTING -p TCP --dport 763 -j DNAT --to-destination 10.1.1.20
iptables -t nat -A PREROUTING -i eth2 -p UDP --dport 1023 -j REDIRECT --to-port 6277
iptables -t nat -A PREROUTING -p UDP --dport 6277 \
-j DNAT --to-destination 10.1.1.20:1023
iptables -t nat -A PREROUTING -j LOG
echo " Done.\n"

This section sets up IP masquerading, natting all packets that
leave eth0 or eth1 to the addresses assigned to eth0 and eth1.
echo " Setting masquerading : "
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
iptables -t nat -A POSTROUTING -o eth1 -j MASQUERADE
echo " Done.\n"

echo -e "\n$name script, v$version done.\n"

```

4. If you are using the OSS filtering and Domino mail delivery model for your system, enter the following lines before the last line in the file:

```

echo " Setting Domino delivery forwarding rules..."
iptables -t nat -A PREROUTING -i eth2 -p TCP --dport 25 \
-j DNAT --to-destination 172.27.72.20
iptables -t nat -A PREROUTING -p TCP --dport 110 \

```

```
-j DNAT --to-destination 172.27.72.20
echo " Done."
```

If you are using the OSS filtering and OSS mail delivery model for your system, enter the following lines before the last line in the file:

```
echo " Setting OSS delivery forwarding rules..."
iptables -t nat -A PREROUTING -p TCP --dport 110 \
-j DNAT --to-destination 10.1.1.20
echo " Done."
```

**Note:** This iptables ruleset blocks all incoming connections unless they are explicitly defined or related to an already existing connection. If you require other services on your network to be accessible from your LAN or the Internet, simply copy the format of the above rules to create new rules. The new rules must allow connections on whatever port the required service uses, and then forward those connections to the required internal server.

Protocols and port numbers recognized by your system can be found in the `/etc/services` file.

5. Press the Esc button to exit editing mode.
6. Type `:wq` to save the file and exit the vi editor.
7. Make the file executable:

```
~> chmod a+x /etc/firewall.start
```
8. Activate all the rules now by running the script:

```
~> ./etc/firewall.start
```
9. Now add the script to your system boot sequence to reset the rules every time the server is restarted:
  - a. Open the `boot.local` file for editing in the vi text editor:

```
~> vi /etc/rc.d/boot.local
```
  - b. Press the Insert key to enter editing mode.
  - c. After the last line in the file, enter the following text:

```
./etc/firewall.start
```
  - d. Exit text entry mode by pressing the Esc key.
  - e. Save the file and exit the vi text editor by typing `:wq`.

The firewall is now configured properly, and the rules have been activated. At system boot, the rules will be activated.

**Note:** For maximum security, any physical Ethernet adapters assigned to the i5/OS should now be unplugged and unassigned. A connection to administer the i5/OS can now be established through the firewall, and via virtual Ethernet connections. Client access should be replaced with the Web-based client access version. Those wishing to connect from the LAN need simply point their clients to the address of the firewall (cell BB2 on the planning worksheet). The firewall will forward HTTP (port 80) connections to the i5/OS partition. Those wishing to connect to the i5/OS from the Internet simply need to point their clients to the Internet routable address assigned to the eServer i5 (cell CB2 on the planning worksheet), on port 443. The firewall will forward secure HTTP (port 443) connections to the i5/OS.

### 3.3 grsecurity kernel patch

The grsecurity patch is an addition to the Linux kernel that provides several well implemented security enhancements to your system. The grsecurity patch institutes network socket control, Mandatory Access Control (MAC), randomization of local and network informational data, /proc restrictions, fine-grained auditing, and many more features that will help to secure your system against attackers. To install the grsecurity kernel patch, follow these steps:

1. You will need to be in the src directory on your system. Switch there:

```
~> cd /usr/src
```

2. Download the 2.6.11.7 kernel source:

```
~> curl -O http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.11.7.tar.bz2
```

3. Extract the kernel source:

```
~> tar xjf linux-2.6.11.7.tar.bz2
```

4. Download the grsecurity patch for the 2.6.11.7 kernel:

```
~> curl -O http://www.grsecurity.net/grsecurity-2.1.5-2.6.11.7-200504111924.patch.gz
```

5. Extract the grsecurity patch:

```
~> gunzip grsecurity-2.1.5-2.6.11.7-200504111924.patch.gz
```

6. Patch the kernel source:

```
~> patch -p0 < ./grsecurity-2.1.5-2.6.11.7-200504111924.patch
```

7. Remove old files:

- a. Switch to the kernel source directory:

```
~> cd linux-2.6.11.7
```

- b. Compile the clean make target:

```
~> make clean
```

8. Configure the kernel:

- a. Load the kernel configuration menu:

```
~> make menuconfig
```

- b. You will be presented with the Linux Kernel v2.6.11.7-grsec Configuration screen, as shown in Figure 3-67 on page 167. Use the arrow keys to select Security options from the list, and press the Enter key.

```

Linux Kernel v2.6.11.7-grsec Configuration

+----- Linux Kernel Configuration -----+
| Arrow keys navigate the menu. <Enter> selects submenus --->.
| Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes,
| <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help, </>
| for Search. Legend: [*] built-in [] excluded <M> module < >
+-----+
| Code maturity level options --->
| General setup --->
| Loadable module support --->
| Platform support --->
| General setup --->
| Device Drivers --->
| File systems --->
| Profiling support --->
| Kernel hacking --->
| Security options --->
+-----(+)-+
| <Select> < Exit > < Help >
+-----+

```

Figure 3-67 grsecurity - Selecting Security options

- c. Select Grsecurity from the list shown in Figure 3-68, and press the Enter key.

```

Linux Kernel v2.6.11.7-grsec Configuration

+----- Security options -----+
| Arrow keys navigate the menu. <Enter> selects submenus --->.
| Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes,
| <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help, </>
| for Search. Legend: [*] built-in [] excluded <M> module < >
+-----+
| Grsecurity --->
| PaX --->
| [] Enable access key retention support (NEW)
| [*] Enable different security models
| [*] Socket and Networking Security Hooks
| <M> Default Linux Capabilities
| <M> Root Plug Support
| < > BSD Secure Levels (NEW)
| [*] NSA SELinux Support
| [*] NSA SELinux boot parameter
+-----(+)-+
| <Select> < Exit > < Help >
+-----+

```

Figure 3-68 grsecurity - Selecting Grsecurity

- d. You are presented with the Grsecurity menu, with only one option, as shown in Figure 3-69 on page 168. Press the Spacebar to activate the Grsecurity option.





```
Linux Kernel v2.6.11.7-grsec Configuration

+-----+
| Do you wish to save your new kernel configuration? |
+-----+
| < Yes > < No > |
+-----+
```

Figure 3-72 grsecurity - Save kernel configuration

9. Compile and install any selected modules:

a. Compile the modules make target:

```
~> make modules
```

b. Compile the modules\_install make target.

```
~> make modules_install
```

10. Compile the kernel source:

```
~> make
```

Kernel compilation can take quite a while. Be patient until compilation is complete.

11. Install the compiled source:

```
~> make install
```

12. Set the new kernel to be booted by default:

a. If you are running SLES9, follow these instructions to set your boot options:

i. Erase the existing lilo.conf file:

```
~> rm /etc/lilo.conf
```

ii. Open lilo.conf for editing in the vi text editor:

```
~> vi /etc/lilo.conf
```

iii. Press the Insert key to enter text editing mode.

iv. Insert the follow lines:

```
default = linux-gre
timeout = 100
boot = /dev/sda1
activate
```

```
image = /boot/vmlinuz
```

```
label = linux-gre
root = /dev/sda3
initrd = /boot/initrd
append = "selinux=0 elevator=cfq"
```

```
image = /boot/vmlinux
label = linux
root = /dev/sda3
initrd = /boot/initrd-2.6.5-7.97-pseries64
append = "selinux=0 elevator=cfq"
```

- v. Press the Esc key to exit editing mode.
- vi. Type :wq to save the file and exit the vi text editor.
- vii. Write the yaboot.conf file:

```
~> lilo
```

- b. If you are running RHEL4, to set your boot options, add the new kernel to yaboot.conf using grubby:

```
grubby --yaboot --add-kernel /boot/KERNELBOOTIMAGE --initrd
/boot/INITRDFILEINITRDFILE --title linux-gre --copy-default --make-default
```

The grsecurity kernel patch is now installed. Upon reboot, the new kernel will be loaded.

## 3.4 Security-Enhanced Linux (SELinux)

SELinux is a kernel security patch developed by the National Security Agency that is now integrated into all kernel releases Version 2.6 or later from the main Linux kernel repository, Kernel.org. SELinux aims for many of the same goals as the grsecurity patch, such as mandatory access controls, in addition to the implementation of new security measures, such as confining super user accounts for additional security.

The SELinux security patch is installed by default with the 2.6.11.7 kernel that was installed when the grsecurity kernel patch was applied. Default security options for SELinux are set automatically, and will protect your system properly without additional configuration.

## 3.5 Snort

This section discusses how to install and configure the Snort network Intrusion Detection System (IDS). Snort analyzes network traffic and matches patterns it finds to a predefined set of rules. When Snort detects malicious activity, it executes a preset action, such as sending an e-mail to a system administrator, alerting them of the problem.

### 3.5.1 Installing libpcap 0.9.0-096

The libpcap library is an OSS packet capturing library that allows for the analyzation of packets. Snort requires the libpcap library for proper network protection. To install libpcap, follow these steps:

1. Open up the PuTTY client.
2. Enter the IP address of the LINUXFW partition from cell CB2 on the planning worksheet in the Host Name (or IP address) field.
3. Set the protocol to SSH.
4. Click the **Open** button to open the connection.

5. Log in to your system using your regular user name and password from cells P2 and Q2 on the planning worksheet.
6. Switch to the /usr/src directory:  
~> cd /usr/src
7. Download the libpcap library:  
~> curl -O http://www.tcpdump.org/release/libpcap-0.8.3.tar.gz
8. Extract the downloaded source:  
~> tar xzf libpcap-0.8.3.tar.gz
9. Change to the libpcap directory:  
~> cd libpcap-0.8.3
10. Configure, compile, and install the extracted source code:  
~> ./configure; make; make install
11. Exit the libpcap directory by:  
~> cd ..

The libpcap libraries has been installed.

### 3.5.2 Installing Perl Compatible Regular Expressions (PCRE) 5.0

The PCRE functions implement regular expression pattern matching in the same way as the Perl 5 programming language. Snort uses the PCRE libraries to function. To install the PCRE libraries, follow these instructions:

1. Download the PCRE 5.0 tarball:  
~> curl -O http://easynews.dl.sourceforge.net/sourceforge/pcre/pcre-5.0.tar.gz
2. Extract the downloaded source:  
~> tar xzf pcre-5.0.tar.gz
3. Change to the PCRE directory:  
~> cd pcre-5.0
4. Configure, compile, and install the PCRE source code:  
~> ./configure; make; make install
5. Exit the PCRE directory:  
~> cd ..

The PCRE libraries have been installed.

### 3.5.3 Installing Snort 2.3.1

Snort checks every packet passing through the firewall for potential dangers to your system. Once the libpcap and PCRE libraries have been installed, follow these instructions to install Snort:

1. Download the Snort 2.1.3 tarball:  
~> curl -O http://www.snort.org/dl/old/snort-2.3.1.tar.gz
2. Extract the source:  
~> tar xzf snort-2.3.1.tar.gz

3. Change to the Snort directory:  

```
~> cd snort-2.3.1
```
4. Configure, compile, and install the PCRE source code:  

```
~> ./configure; make; make install
```
5. Exit the PCRE directory:  

```
~> cd ..
```
6. Make a snort directory in /etc, and copy the snort configuration file to it:  

```
~> mkdir /etc/snort
~> cp /etc/tools/snort-2.3.1/etc/snort.conf /etc/snort/snort.conf
```
7. Make a rules directory in /etc/snort and copy all the default Snort rules to it:  

```
~> mkdir /etc/snort/rules
~> cp /etc/tools/snort-2.3.1/rules/* /etc/snort/rules
```
8. Make a snort log directory:  

```
~> mkdir /var/log/snort
```

Snort has been installed and is ready for configuration.

### 3.5.4 Configuring Snort

To configure Snort:

1. Open the snort.conf file for editing in the vi text editor:

```
~> vi /etc/snort/snort.conf
```

2. Press the Insert key to enter editing mode.

3. Leave the following lines intact:

```
var HOME_NET any
```

```
var EXTERNAL_NET any
```

This forces snort to check all packets that pass through all interfaces for both internal and external style attacks. This may seem wasteful; however, a spoofed internal IP address could easily bypass an IDS even if the attack was originating from an external source. Setting both variables to all network address is also useful if your internal networks span multiple subnets, in which case defining each one of them can be a painstaking task.

4. Set your SMTP server. Change the follow line:

```
var SMTP_SERVERS $HOME_NET
```

to read:

```
var SMTP_SERVERS 10.1.1.20
```

5. Set the path to the rule files. Change the line:

```
var RULE_PATH ../rules
```

to read:

```
var RULE_PATH /etc/snort/rules
```

6. Set the stream4 preprocessor options to detect port scans. Change the following line:

```
preprocessor stream4: disable_evasion_alerts
```

to read:

```
preprocessor stream4: detect_scans disable_evasion_alerts
```

This tells Snort to disassemble an incoming IPv4 stream, and to check that disassembled stream for evidence of port scans or detection evasion.

7. Set the stream4\_reassemble preprocessor options to reassemble both client and server TCP streams on all default ports. Change the following line:

```
preprocessor stream4_reassemble
```

to read:

```
preprocessor stream4_reassemble: both ports "default"
```

8. Define the system log alert plugin. Uncomment the following line:

```
output alert_syslog: LOG_AUTH LOG_ALERT
```

9. Enable attack rules. Uncomment the following lines:

```
include $RULE_PATH/web-attacks.rules
include $RULE_PATH/backdoor.rules
include $RULE_PATH/policy.rules
include $RULE_PATH/info.rules
include $RULE_PATH/icmp-info.rules
```

**Note:** Not all of the enabled rules will be necessary for all systems. However, having Snort check traffic for attacks that will have no effect on your system does not cause any damage. It is better to log harmless attacks than let harmful attacks through.

However, if your system receives a high volume of messages or connection requests, the processing power required to process every request may become demanding. If this is the case, you may consider slimming down your ruleset to only check for attacks that would have an effect on your system.

10. Press the Esc key to exit editing mode.

11. Save the file by typing :wq.

12. Run Snort in daemon mode:

```
~> /usr/local/bin/snort -c /etc/snort/snort.conf -D
```

13. Add Snort to your system startup:

- a. Open the boot.local file for editing in the vi text editor:

```
~> vi /etc/rc.d/boot.local
```

- b. Press the Insert key to enter editing mode.

- c. After the last line in the file, enter the following text:

```
/usr/local/bin/snort -c /etc/snort/snort.conf -D
```

- d. Exit text entry mode by pressing the Esc key.

- e. Save the file and exit the vi text editor by typing :wq.

Snort is now installed and running. Snort will load at system boot, and write out alerts to the system logs in addition to multiple log files located in /var/log/snort.

**Important:** Snort alerts are logged, but not processed for any further action. It is within the responsibilities of the system administrator to check the system and Snort-generated logs for dangerous alerts that require attention. Using additional open source software, such as the popular Snortlog log interpretation system, it is possible to have Snort send e-mails or pages to the system administrator if an alert is generated. However, this kind of advanced setup is outside of the scope of this redpaper.

## 3.6 Rootkit hunter

A rootkit is a set of tools used by an attacker after they have gained access to your system. A rootkit is used by attackers to maintain access to a cracked system, even after the initial security hole through which they gained access has been patched. A rootkit hunter is a tool that will scan your system for known rootkit tools, files, and common alterations made to files by rootkits. To install the rootkit hunter software on your server, follow these steps:

1. Ensure you are in the directory created for the security tools:  

```
~> cd /etc/tools
```
2. Download the rootkit hunter tarball:  

```
~> curl -O http://downloads.rootkit.nl/rkhunter-1.2.5.tar.gz
```
3. Extract the tarball:  

```
~> tar xzf rkhunter-1.2.5.targ.gz
```
4. Switch to the rkhunter directory:  

```
~> cd rkhunter
```
5. Run the installer script:  

```
~> ./installer.sh
```
6. Set the rootkit hunter to scan the system once weekly, at midnight every Sunday:
  - a. Open the crontab file for editing in the vi text editor:  

```
~> vi /etc/crontab
```
  - b. Enter text entry mode by pressing the Insert key.
  - c. Insert the following line after the last line in the file:  

```
00 0 * * 7 root /etc/tools/rkhunter/files/rkhunter -c --createlogfile
--cronjob
```
  - d. Exit text entry mode by pressing the Esc key.
  - e. Save the file by typing :wq.

The rootkit hunter will run every Sunday at midnight, and log the results of all scans to `/var/log/rkhunter`.

**Important:** Any rootkit that the rootkit hunter finds will be logged to `/var/log/rkhunter`. It is within the responsibilities of the system administrator to check the logs weekly and take appropriate action if any rootkits have been found.

Removing a rootkit from your system can be an extremely complicated and time-consuming process. In the event that a rootkit is detected, we recommend that you immediately remove your system from any networks it is on. Back up all data and logs, and perform a complete re-install of all operating systems and software. Restore your backed up data. Analyze the backed up logs to determine how your system was compromised, and take appropriate action to fix the security hole.

Archived

## E-mail Security tools installation and configuration

This chapter outlines the installation and configuration procedures for mail transfer agent's (MTA's) Postfix and qmail and their security tools Clam AntiVirus and SpamAssassin. For proper functionality of Postfix and qmail, Clam AntiVirus and SpamAssassin are installed on the LINUXST partition.

In this chapter we discuss two MTAs, Postfix and qmail. Also, we integrate them with mail security tools. You can choose either Postfix or qmail. By choosing one MTA you can skip the sections specific to other MTAs. In our opinion, Postfix is much easier than qmail in terms of installation and configuration, as it involves less steps than qmail.

**Attention:** This chapter covers installation of many open source packages. In our examples, we have given direct links to download the packages. If any of the links do not work, it means either the package is replaced by a new version or the mirror site is not available from your location for the time being. If you have any problem downloading those packages using the given URLs, please visit the relevant Web site, as listed in the "Packages: Links to download" on page 260, to get the valid link for downloading. Then replace the URL with our example.

## 4.1 Postfix

Postfix is an alternative to Sendmail. Sendmail is considered to be the de facto mail server for UNIX/Linux systems. It includes many features needed in a mail system and has been in circulation for years, and as such provides excellent stability.

The only problem with Sendmail is the configuration of the service. Most Linux services use a common configuration file format, commonly known as rc files. The rc files use familiar variable=value pairs that are usually self explanatory.

Sendmail, on the other hand, uses cryptic variable/value pairs. Having a system that is difficult to configure and administer can lead to the administrator overlooking or wrongly setting configuration options.

Postfix was written by Wietse Venema while he was a resident at an IBM research facility. In Wietse's own words: "Postfix is my attempt to provide an alternative to the widely used Sendmail program. Postfix attempts to be fast, easy to administer, and hopefully secure, while at the same time being sendmail-compatible enough to not upset your users."

You can find more details on Postfix on its home page at:

<http://www.postfix.org>

### 4.1.1 Preparing to install Postfix

The installation of Postfix is not as straightforward as other applications. If Sendmail is installed in a system, remove it before you install Postfix. You are going to run Postfix instead of Sendmail. Running both mail servers in the same system will produce unpredictable results and will almost definitely prohibit one of the servers from running.

For SLES9 users, if you have installed SLES9 with the minimum installation option, by default, you have Postfix installed. This means you do not have Sendmail.

For RHEL4 users, since the RHEL4 minimum installation includes both Sendmail and Postfix with a Mail Transfer Switcher Agent, just use the Mail Transfer Switcher Agent to switch from Sendmail to Postfix. The Mail Transfer Switcher Agent does the necessary changes to run the alternate MTA.

In any other case, if you happen to have a running Sendmail, you have to remove it before installing Postfix. This section describes the steps to remove Sendmail.

1. Log in to the LINUXST partition following the steps below.
  - a. Open a virtual console using PuTTY. Enter your LINUXFW IP address, set your protocol to SSH, and set your port value to 22. This opens a screen prompting you to login. Type your user name and password. On successful login the screen looks as shown in Figure 4-1 on page 179.

**Attention:** Since there is no direct login, to access the LINUXST partition you must log in to the LINUXFW partition first. From there you can access the LINUXST partition.

```
login as: smail
Using keyboard-interactive authentication.
Password:
Last login: Tue May 3 11:29:51 2005 from p78fczpf.rchland.ibm.com
smail@LINUXFW:~>
```

*Figure 4-1 Logging in to LINUXFW partition*

- b. To log in to the LINUXST partition type at the command prompt:  
# ssh root@<IP address of LINUXST partition>
- c. When you log in from LINUXFW to LINUXST for the very first time, ssh will ask whether to add the host in the list of known hosts. Type yes and type the root password of the LINUXST partition. On successful login, the screen looks as shown in the Figure 4-2 on page 180.

```

login as: smail
Using keyboard-interactive authentication.
Password:
Last login: Tue May 3 12:55:42 2005 from p78fczpf.rchland.ibm.com
smail@LINUXFW:~> ssh root@10.1.1.20
The authenticity of host '10.1.1.20 (10.1.1.20)' can't be established.
RSA key fingerprint is 19:e1:58:ea:70:47:ff:96:cb:57:4b:95:38:41:8d:a3.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.1.1.20' (RSA) to the list of known hosts.
Password:
Last login: Tue May 3 10:55:00 2005 from 10.1.1.10
LINUXST:~ #

```

Figure 4-2 Logging in to LINUXST partition

## 2. Remove Sendmail.

First you have to check whether your Sendmail installation is an RPM installation or a source installation. Type the following rpm query command:

```
rpm -qa | grep sendmail
```

If Sendmail is installed as an RPM, you will see output similar to this:

```
sendmail-x.x.x.x
sendmail-doc-x.x.x.x
sendmail-devel-x.x.x.x
sendmail-cf-x.x.x.x
```

If it is installed as an RPM, you can uninstall Sendmail quite easily by running the following commands:

```
/etc/rc.d/init.d/sendmail stop
rpm -e --nodeps sendmail
rpm -e --nodeps sendmail-doc
rpm -e --nodeps sendmail-devel
rpm -e --nodeps sendmail-cf
```

However, if the RPM query produces no output, you most likely have a source install of Sendmail. In that case, type the following commands to stop and disable Sendmail.

```
/etc/rc.d/init.d/sendmail stop
mv /usr/lib/sendmail /usr/lib/sendmail.old
mv /usr/sbin/sendmail /usr/sbin/sendmail.old
chmod 0 /usr/lib/sendmail.old /usr/sbin/sendmail.old
```

Since Sendmail is removed, you can start updating Postfix now.

## 4.1.2 Updating Postfix

The Postfix SMTP server must be upgraded to ensure proper functionality by installing the latest version of Postfix as explained in the following steps:

1. With the virtual console still open, stop the currently running Postfix process by executing the command:

```
postfix stop
```

2. Ensure you are in the `/usr/local/src` directory by executing the command:

```
cd /usr/local/src
```

3. Download the Postfix source code by executing the command:

```
curl -O http://postfix.energybeam.com/source/official/postfix-2.2.2.tar.gz
```

**Note:** To download files from the URLs we can use either `wget` or `curl`. Since the minimal installation of SLES9 does not include `wget`, we are using `curl` to download the files instead. The option `-O` in `curl` specifies that the file name to be written is same as in the source. If you prefer to use `wget`, you can install the `wget` rpm from media #2 for SLES9. In case of RHEL4, both `wget` and `curl` are available by default.

4. Extract the downloaded source by executing the command:

```
tar xzf postfix-2.2.2.tar.gz
```

5. Change to the Postfix directory by executing the command:

```
cd postfix-2.2.2
```

6. Compile the Postfix by executing the command:

```
make
```

7. Install the Postfix by executing the command:

```
make install
```

**Note:** Please note that you can do `make install` only after the successful completion of `make`. The successful completion means the command should end with *no* errors. The same rule applies for all packages we are going to install from source files similar to this one. It is understood that wherever we use the commands `configure`, `make`, and `make install` they end with no errors. If there are any errors, they should be fixed before proceeding to next step. The common errors would be on file permissions and dependencies. We have included all necessary dependency programs for our setup. Make sure you have not missed any such packages.

Figure 4-3 on page 182 shows the steps involved in Postfix installation.

```

LINUXST:~ # postfix stop
postfix/postfix-script: stopping the Postfix mail system
LINUXST:~ # cd /usr/local/src/
LINUXST:/usr/local/src # curl -O
http://postfix.energybeam.com/source/official/postfix-2.2.2.tar.gz
 % Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 0k 0 0:00:02 0:00:02 0:00:00 0k
LINUXST:/usr/local/src # tar xzf postfix-2.2.2.tar.gz
LINUXST:/usr/local/src # cd postfix-2.2.2/
LINUXST:/usr/local/src/postfix-2.2.2 # make

<part of the output is skipped here>

gcc -Wmissing-prototypes -Wformat -DHAS_PCRE -I/usr/local/include -g -O -I.
-I../include -DLINUX2 -o tlmgr tlmgr.o ../lib/libmaster.a ../lib/libtls.a
../lib/libglobal.a ../lib/libutil.a -L/usr/local/lib -lpcrc -ldb -lnsl -lresolv
cp tlmgr ../libexec
LINUXST:/usr/local/src/postfix-2.2.2 # make install

<output is skipped here>

```

Figure 4-3 Installing Postfix

8. The Postfix installation program will now ask you a series of questions like `install_root`, `tempdir`, `config_directory`, and so on. The default answers are all fine for installing Postfix properly on your system. Simply press Enter for each question you are prompted with.

Figure 4-4 on page 183 shows a few such questions with the default answers we proceed with.

```

<the initial output of make install is skipped here>

/bin/sh postfix-install

Warning: if you use this script to install Postfix locally,
this script will replace existing sendmail or Postfix programs.
Make backups if you want to be able to recover.

Before installing files, this script prompts you for some definitions.
Most definitions will be remembered, so you have to specify them
only once. All definitions should have a reasonable default value.

Please specify the prefix for installed file names. Specify this ONLY
if you are building ready-to-install packages for distribution to other
machines.
install_root: [/]

Please specify a directory for scratch files while installing Postfix. You
must have write permission in this directory.
tempdir: [/usr/local/src/postfix-2.2.2]

Please specify the final destination directory for installed Postfix
configuration files.
config_directory: [/etc/postfix]

<remaining part of output is skipped here>

```

Figure 4-4 Installing Postfix (continued)

### 4.1.3 Postfix configuration files

Once you have installed Postfix you can start configuring it. Postfix's configuration files are placed in the `/etc/postfix` directory. The configuration files and their purposes are listed below.

|                  |                                                                                                                                                                                                                                                          |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>access</b>    | This configuration file dictates which hosts are allowed to connect to Postfix.                                                                                                                                                                          |
| <b>aliases</b>   | A configurable list required by the mail protocol.                                                                                                                                                                                                       |
| <b>canonical</b> | This defines address remapping rules, both local and remote.                                                                                                                                                                                             |
| <b>main.cf</b>   | This file is the global configuration file for Postfix.                                                                                                                                                                                                  |
| <b>master.cf</b> | This defines the master daemon configuration that controls Postfix daemons for incoming and outgoing mail.                                                                                                                                               |
| <b>relocated</b> | This defines members of the organization that are no longer using the mail system. Postfix will notify senders that the recipient is no longer at the address and can also tell the sender of a new e-mail address where the recipient can be contacted. |
| <b>transport</b> | This maps e-mail addresses to relay hosts.                                                                                                                                                                                                               |

Among these files, the two most important configuration files are `/etc/postfix/main.cf` and `/etc/postfix/master.cf`

## 4.1.4 Configuring Postfix

Postfix uses two configuration files (located at `/etc/postfix/main.cf` and `/etc/postfix/master.cf`) to define its behavior. To configure Postfix to function as a standalone SMTP server with no virus or spam scanning, only the `main.cf` file needs to be edited.

1. Edit the Postfix config file:

a. To start editing the configuration file, at the command line type:

```
vi /etc/postfix/main.cf
```

b. Make sure of the existence of an uncommented line that reads:

```
queue_directory = /var/spool/postfix
```

which sets the location of the Postfix top-level queue directory.

**Note:** To set up your SMTP server, you must have your own fully qualified domain name. In this example, the domain name that will be used is `secureyouremails.com`. Substitute your own domain name wherever `secureyouremails.com` is used. It is customary to set the host name of your SMTP server to `mail.yourhostname.com`; however, this is not required. Ensure that the MX records for your domain's DNS point to your mail server.

c. Uncomment the line that reads:

```
myhostname = host.domain.tld
```

and change the line so that it reads,

```
myhostname = mail.secureyouremails.com
```

d. Uncomment the line that reads:

```
mydomain = domain.tld
```

and change the value of the `mydomain` parameter to read `secureyouremails.com` like:

```
mydomain = secureyouremails.com
```

e. Uncomment the line that reads:

```
myorigin = $mydomain
```

The `myorigin` parameter specifies the domain that appears in mail that is posted on this machine. The default is to use the local machine name, `$myhostname`, which defaults to the name of the machine.

f. Uncomment the line that reads:

```
inet_interfaces = all
```

The `inet_interfaces` parameter specifies all network interface addresses that the Postfix system should listen on.

g. Uncomment the line that reads:

```
relayhost = [an.ip.add.ress]
```

and change the line so that it reads:

```
relayhost = MAILSERVERIP
```

**Attention:** Please note that MAILSERVERIP must be replaced with the IP address of the Mail Server (SMTP). If you are using Lotus Domino on i5, this IP address is the same as the virtual IP of the LINUXFW partition. The iptable rule on the LINUXFW partition will route the messages to Domino in turn. If you are not using Lotus Domino then this IP address is the same as the virtual IP of the LINUXST partition.

- h. Uncomment the line that reads:

```
local_recipient_maps = (i.e. empty)
```

and change this line so that it reads:

```
local_recipient_maps =
```

**Note:** The `local_recipient_maps` parameter specifies optional lookup tables with all names or addresses of users that are local with respect to `$mydestination`, `$inet_interfaces`, or `$proxy_interfaces`. If this parameter is defined, then the SMTP server will reject mail for unknown local users. This parameter is defined by default. To turn off local recipient checking in the SMTP server, specify `local_recipient_maps = (that is, empty)`

- i. At the end of the file, the Postfix installation program has inserted default settings, many of which are no longer necessary. The following lines need to be commented out:

```
mail_spool_directory = /var/mail
myhostname = linux.site
inet_interfaces = 127.0.0.1 ::1
mydestination = $myhostname, localhost.$mydomain
relayhost =
```

- j. Save the file and exit the editor.

2. Create the spool directory. At the command line type:

```
mkdir /var/spool/postfix
```

3. Now start Postfix. At the command line type:

```
postfix start
```

Postfix is now configured to receive mail and forward it to your mail server. Postfix writes any output to a log file located at `/var/log/mail`. The log file may be checked to obtain information about the status of Postfix.

## 4.2 qmail

qmail is another Sendmail alternative for UNIX-like operating systems. qmail is more secure, efficient, reliable, and simple. When compared to any other equivalently featured MTA, qmail is smaller.

One of the known strengths of qmail is its mailing list management. qmail lets each user handle his own mailing lists through easy setup. It also automatically prevents mailing list loops, even across hosts. If you are looking for a more secure and efficient mail transfer agent then qmail is the right choice.

qmail's modular, lightweight design and sensible queue management make it the fastest MTA.

**Note:** It is just *qmail* not Qmail, QMail, or qMail.

You can find more details on qmail from its official site at:

<http://cr.yp.to>

## 4.2.1 Overview of qmail installation

The qmail installation is not as simple as other MTAs. Apart from installing the qmail itself there are many other packages to be installed.

Since qmail follows the classic UNIX philosophy that each tool should perform a single, well-defined function, and complex functions should be built by connecting a series of simple tools into a pipeline, qmail itself does not do everything.

Here, then, are some of the most popular add-ons written for qmail. Of course, many standard UNIX utilities can also be plugged into qmail.

- ▶ dot-forward: A Sendmail *.forward* file compatibility add-on
- ▶ fastforward: A Sendmail alias database compatibility add-on
- ▶ ucspi-tcp: An inetd replacement
- ▶ daemontools: A set of tools for managing daemons and their logs
- ▶ qmailanalog: A set of qmail log file analysis tools
- ▶ ezmlm: A mailing list manager for qmail

**Note:** A list of add-on packages is available at qmail's official site:

<http://cr.yp.to/software.html>

Also, there are many third-party tools available for qmail. You can install them depending on your requirements. For example the tool qmailanalog is to help analyze qmail's activity record. It gives statistics like how many messages are being delivered, how many recipients, how many successes and failures, and so on. We do not discuss all such tools here. Only the minimum required tools are to be installed and discussed.

Our installation of qmail includes the following packages.

- ▶ netqmail: This consists of the official qmail 1.03 and a patch that fixes a handful of bugs, deficiencies, and incompatibilities.
- ▶ ucspi-tcp: This is for the inetd replacement.
- ▶ daemontools: This consists of a set of tools for managing daemons and their logs.
- ▶ fastforward: This add-on provides the Sendmail alias database compatibility.
- ▶ maildrop: This is a mail filter/mail delivery agent required for qmail-scanner.
- ▶ qmailscanner: This is to integrate Clam Antivirus and SpamAssassin into our qmail server's mail queue.

## 4.2.2 Preparing to install qmail

Before installing qmail, other MTAs like Sendmail or Postfix, if any are installed, should be removed. In this section we discuss removing Sendmail and Postfix.

- ▶ Removing Sendmail

First you have to check whether your Sendmail installation is an RPM installation or a source installation. Type the following rpm query command.

```
rpm -qa | grep sendmail
```

If Sendmail is installed as an RPM, you will see output similar to this:

```
sendmail-x.x.x.x
sendmail-doc-x.x.x.x
sendmail-devel-x.x.x.x
sendmail-cf-x.x.x.x
```

If it is installed as an RPM, you can uninstall Sendmail quite easily by running the following commands:

```
/etc/rc.d/init.d/sendmail stop
rpm -e --nodeps sendmail
rpm -e --nodeps sendmail-doc
rpm -e --nodeps sendmail-devel
rpm -e --nodeps sendmail-cf
```

However, if the RPM query produces no output, you most likely have a source install of Sendmail. In that case type the following commands to stop and disable Sendmail.

```
/etc/rc.d/init.d/sendmail stop
mv /usr/lib/sendmail /usr/lib/sendmail.old
mv /usr/sbin/sendmail /usr/sbin/sendmail.old
chmod 0 /usr/lib/sendmail.old /usr/sbin/sendmail.old
```

#### ► Removing Postfix

Type the following rpm query to obtain the installed Postfix's rpm.

```
rpm -qa | grep postfix
```

If it is installed, you will get output similar to this:

```
postfix-x.x-x
```

Remove the Postfix by executing the commands below.

```
/etc/init.d/postfix stop
rpm -e --nodeps postfix
```

Figure 4-5 on page 188 shows the steps to remove Postfix.

```

LINUXST:~ # rpm -qa | grep postfix
postfix-2.1.1-1.4
LINUXST:~ # /etc/init.d/postfix stop
Shutting down mail service (Postfix) done
LINUXST:~ # rpm -e --nodeps postfix-2.1.1-1.4
warning: /etc/postfix/virtual saved as /etc/postfix/virtual.rpmsave
warning: /etc/postfix/sender_canonical saved as /etc/postfix/sender_canonical.rpmsave
warning: /etc/postfix/main.cf saved as /etc/postfix/main.cf.rpmsave
LINUXST:~ #

```

Figure 4-5 Removing Postfix

Since qmail's installation program creates the subdirectories as they are needed, you have to create the qmail home directory. The home directory includes the configuration files, the queue, and the actual qmail binaries. At the command prompt type:

```
mkdir /var/qmail
```

### Creating qmail users and groups

qmail requires the existence of a list of users and groups before doing the installation. To create those users and groups type the series of commands below.

```

groupadd nofiles
useradd -g nofiles -d /var/qmail/alias -s /sbin/nologin -p '*' alias
useradd -g nofiles -d /var/qmail -s /sbin/nologin -p '*' qmaild
useradd -g nofiles -d /var/qmail -s /sbin/nologin -p '*' qmail1
useradd -g nofiles -d /var/qmail -s /sbin/nologin -p '*' qmailp
groupadd qmail
useradd -g qmail -d /var/qmail -s /sbin/nologin -p '*' qmailq
useradd -g qmail -d /var/qmail -s /sbin/nologin -p '*' qmailr
useradd -g qmail -d /var/qmail -s /sbin/nologin -p '*' qmails

```

**Note:** If these users and groups are not available, you cannot install qmail.

## 4.2.3 Installing qmail

As discussed earlier, we are going to install the following packages, which ensure the required qmail function.

- ▶ netqmail-1.05
- ▶ uscp-tcp-0.88
- ▶ daemontools-0.76
- ▶ fastforward-0.51

- ▶ maildrop-1.8.1
- ▶ qmailscanner-1.25

To install:

#### 1. Installing netqmail

To install netqmail-1.05, follow the steps below:

- a. Change the present working directory to `/usr/local/src` by executing the command:

```
cd /usr/local/src
```

- b. Download netqmail-1.05 from <http://www.qmail.org>. At the command prompt type:

```
curl -O http://www.qmail.org/netqmail-1.05.tar.gz
```

- c. Extract the source by executing the command:

```
tar xzf netqmail-1.05.tar.gz
```

- d. Change to the directory netqmail-1.05 by executing the command:

```
cd netqmail-1.05
```

- e. As mentioned earlier, the package netqmail includes the official qmail and the necessary patches for qmail and other tools. To apply the patches to qmail and make a new source directory, run the script `collate.sh` by executing the command:

```
./collate.sh
```

Now the screen may like Figure 4-6.

```
LINUXST:~ # cd /usr/local/src
LINUXST:/usr/local/src # curl -O http://www.qmail.org/netqmail-1.05.tar.gz
 % Total % Received % Xferd Average Speed Time Curr.
 % Total % Received % Xferd Average Speed Time Curr.
100 0k 100 0k 0 0 0k 0 0:00:02 0:00:02 0:00:00 0k
LINUXST:/usr/local/src # tar xzf netqmail-1.05.tar.gz
LINUXST:/usr/local/src # cd netqmail-1.05/
LINUXST:/usr/local/src/netqmail-1.05 # ./collate.sh
```

You should see 7 lines of text below. If you see anything else, then something might be wrong.

```
[1] Extracting qmail-1.03...
[2] Patching qmail-1.03 into netqmail-1.05. Look for errors below:
24
[4] The previous line should say 24 if you used GNU patch.
[5] Renaming qmail-1.03 to netqmail-1.05...
[6] Continue installing qmail using the instructions found at:
[7] http://www.lifewithqmail.org/lwq.html#installation
LINUXST:/usr/local/src/netqmail-1.05 #
```

Figure 4-6 Installing netqmail

- f. Now there is a directory called netqmail-1.0.5 created under the directory `/usr/local/src/netqmail-1.0.5`. Change to the directory by executing the command:

```
cd /usr/local/src/netqmail-1.0.5/netqmail-1.0.5
```

- g. To create the build and install, execute the following commands:

```
make
make setup check
```

- h. Once build is completed, do post install configuration by running the config script by executing the command:

```
./config-fast mail.secureyouremails.com
```

**Note:** To set up your SMTP server, you must have your own fully qualified domain name. In this example, the domain name that will be used is secureyouremails.com. Please substitute your own domain name wherever secureyouremails.com is used. It is customary to set the host name of your SMTP server to mail.yourhostname.com; however, this is not required. Please ensure that the MX records for your domain's DNS point to your mail server.

Now the screen may appear similar to Figure 4-7.

```
LINUXST:/usr/local/src/netqmail-1.05/netqmail-1.05 # ./config-fast
mail.secureyouremails.com
Your fully qualified host name is mail.secureyouremails.com.
Putting mail.secureyouremails.com into control/me...
Putting secureyouremails.com into control/defaultdomain...
Putting secureyouremails.com into control/plusdomain...
Putting mail.secureyouremails.com into control/locals...
Putting mail.secureyouremails.com into control/rcpthosts...
Now qmail will refuse to accept SMTP messages except to mail.secureyouremails.com.
Make sure to change rcpthosts if you add hosts to locals or virtualdomains!
LINUXST:/usr/local/src/netqmail-1.05/netqmail-1.05 #
```

Figure 4-7 Installing netqmail (continued)

**Note:** qmail requires a working DNS. Before running the config scripts make sure your DNS is set up and running.

- i. That is it. netqmail is installed. Exit the netqmail directory by executing the command:

```
cd ../../
```

## 2. Installing ucspi-tcp

To install ucspi-tcp follow the steps below.

- a. Change the present working directory to /usr/local/src by executing the command:

```
cd /usr/local/src
```

- b. Download ucspi-tcp-0.88 from the gmail's official site <http://cr.yip.to>. At the command prompt type:
 

```
curl -O http://cr.yip.to/ucspi-tcp/ucspi-tcp-0.88.tar.gz
```
- c. Extract the source by executing the command:
 

```
tar xzf ucspi-tcp-0.88.tar.gz
```
- d. Change to the directory ucspi-tcp-0.88 by executing the command:
 

```
cd ucspi-tcp-0.88
```
- e. As mentioned earlier, the package netqmail has necessary patches for other qmail tools, too. To apply the patches to ucspi-tcp-0.88 type the command:
 

```
patch < /usr/local/src/netqmail-1.05/other-patches/ucspi-tcp-0.88.errno.patch
```
- f. Now make the build and install them by executing the commands:
 

```
make
make setup check
```
- g. That is it. ucspi-tcp is installed. Exit the directory by executing the command:
 

```
cd ..
```

Figure 4-8 shows the installation of ucspi-tcp.

```
LINUXST:~# cd /usr/local/src
LINUXST:/usr/local/src # curl -O http://cr.yip.to/ucspi-tcp/ucspi-tcp-0.88.tar.gz
% Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 53019 100 53019 0 0 125k 0 0:00:02 0:00:02 0:00:00 449k
LINUXST:/usr/local/src # tar xzf ucspi-tcp-0.88.tar.gz
LINUXST:/usr/local/src # cd ucspi-tcp-0.88
LINUXST:/usr/local/src/ucspi-tcp-0.88 # patch <
/usr/local/src/netqmail-1.05/other-patches/ucspi-tcp-0.88.errno.patch
patching file error.h
LINUXST:/usr/local/src/ucspi-tcp-0.88 # make

<the output is skipped here>

LINUXST:/usr/local/src/ucspi-tcp-0.88 # make setup check
./install
./instcheck
LINUXST:/usr/local/src/ucspi-tcp-0.88 # cd ..
LINUXST:/usr/local/src #
```

Figure 4-8 Installing ucspi-tcp

### 3. Installing daemontools-0.76

Installing daemontools is somewhat different than installing other packages. Since it has relative paths in the code, it should be downloaded and installed from the directory /package. The following steps explain how to install daemontools-0.76.

- a. Create a directory called package at the root by executing the command:
 

```
mkdir /package
```
- b. Change the present working directory to /package by executing the command:
 

```
cd /package
```
- c. Download daemontools-0.76 from the qmail's official site <http://cr.yp.to>. At the command prompt type:
 

```
curl -O http://cr.yp.to/daemontools/daemontools-0.76.tar.gz
```
- d. Extract the source by executing the command:
 

```
tar xzf daemontools-0.76.tar.gz
```
- e. Change to the daemontools's source directory by executing the command:
 

```
cd admin/daemontools-0.76/src
```
- f. Apply the patches to daemontools-0.76 by executing the command:
 

```
patch < /usr/local/src/netqmail-1.05/other-patches/daemontools-0.76.errno.patch
```
- g. Exit the source directory by executing the command:
 

```
cd ..
```
- h. Now install the daemontools by executing the command:
 

```
package/install
```
- i. That is it. daemontools is installed. Exit the directory by executing the command:
 

```
cd ../../
```

Figure 4-9 shows all the steps involved in installing daemontools-0.76.

```

LINUXST:/usr/local/src # mkdir /package
LINUXST:/usr/local/src # cd /package
LINUXST:/package # curl -O http://cr.yp.to/daemontools/daemontools-0.76.tar.gz
 % Total % Received % Xferd Average Speed Time Time Curr.
 Dload Upload Total Current Left Speed
100 36975 100 36975 0 0 90847 0 0:00:02 0:00:02 0:00:00 319k
LINUXST:/package # tar xzf daemontools-0.76.tar.gz
LINUXST:/package # cd admin/daemontools-0.76/src
LINUXST:/package/admin/daemontools-0.76/src # patch <
/usr/local/src/netqmail-1.05/other-patches/daemontools-0.76.errno.patch
patching file error.h
LINUXST:/package/admin/daemontools-0.76/src # cd ..
LINUXST:/package/admin/daemontools-0.76 # package/install

<some output is skipped here>

Creating /service...
Adding svscanboot to inittab...
init should start svscan now.
LINUXST:/package/admin/daemontools-0.76 # cd ../../
LINUXST:/package #

```

Figure 4-9 Installing daemontools

#### 4. Installing fastforward-0.51

The fastforward add-on handles gmail forwarding according to a cdb database. It can create forwarding databases from a sendmail-style /etc/aliases. The installation is similar to other packages. Follow the quick steps given below to install fastforward.

- a. Change the present working directory to /usr/local/src by executing the command:

```
cd /usr/local/src
```

- b. Download fastforward-0.51 from the qmail's official site, <http://cr.yip.to>. At the command prompt type:

```
curl -O http://cr.yip.to/software/fastforward-0.51.tar.gz
```

- c. Extract the source by executing the command:

```
tar xzf fastforward-0.51.tar.gz
```

- d. Change to the directory fastforward-0.51 by executing the command:

```
cd fastforward-0.51
```

- e. Now make the build and install them by executing the command:

```
make
make setup check
```

- f. That is it. fastforward is installed. Exit the directory by executing the command:

```
cd ..
```

Figure 4-10 shows the installation steps we used.

```
LINUXST:/package # cd /usr/local/src
LINUXST:/usr/local/src # curl -O http://cr.yip.to/software/fastforward-0.51.tar.gz
% Total % Received % Xferd Average Speed Time Curr.
Dload Upload Total Current Left Speed
100 40659 100 40659 0 0 10231 0 0:00:02 0:00:02 0:00:00 461k
LINUXST:/usr/local/src # tar xzf fastforward-0.51.tar.gz
LINUXST:/usr/local/src # cd fastforward-0.51
LINUXST:/usr/local/src/fastforward-0.51 # make

<the output is skipped here>

LINUXST:/usr/local/src/fastforward-0.51 # make setup check
./install
./instcheck
LINUXST:/usr/local/src/fastforward-0.51 # cd ..
LINUXST:/usr/local/src #
```

Figure 4-10 Installing fastforward

## 5. Installing maildrop-1.8.1

Maildrop is required for qmail-scanner. Actually, the whole package of maildrop is not required; only the package reformime, which is available in maildrop, is needed for qmail-scanner.

qmail-scanner needs reformime to extract the MIME attachments. The following steps describe the installation of maildrop-1.8.1.

- a. Change the present working directory to /usr/local/src by executing the command:

```
cd /usr/local/src
```

- b. Download maildrop-1.8.1 from sourceforge.net. At the command prompt type:

```
curl -O
http://easynews.dl.sourceforge.net/sourceforge/courier/maildrop-1.8.1.tar.bz2
```

- c. Extract the source by executing the command:

```
tar xjf maildrop-1.8.1.tar.bz2
```

- d. Change to the directory maildrop-1.8.1 by executing the command:

```
cd maildrop-1.8.1
```

- e. Now configure maildrop with options. At the command prompt type:

```
./configure --prefix=/usr/local --exec-prefix=/usr/local --enable-maildirquota
```

- f. Now make the build and install them by executing the commands one by one.

```
make
make install-strip
make install-man
```

- g. That is it. maildrop is installed. Exit the directory by executing the command:

```
cd ..
```

Figure 4-11 on page 195 shows the installation steps we used.

```

LINUXST:~ # cd /usr/local/src/
LINUXST:/usr/local/src # curl -O
http://easynews.dl.sourceforge.net/sourceforge/courier/maildrop-1.8.1.tar.bz2
 % Total % Received % Xferd Average Speed Time Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 0k 0 0:00:06 0:00:06 0:00:00 0k
LINUXST:/usr/local/src # tar xjf maildrop-1.8.1.tar.bz2
LINUXST:/usr/local/src # cd maildrop-1.8.1
LINUXST:/usr/local/src/maildrop-1.8.1 # ./configure --prefix=/usr/local
--exec-prefix=/usr/local --enable-maildirquota

<output is skipped here>

LINUXST:/usr/local/src/maildrop-1.8.1 # make

<output is skipped here>

LINUXST:/usr/local/src/maildrop-1.8.1 # make install-strip

<output is skipped here>

LINUXST:/usr/local/src/maildrop-1.8.1 # make install-man

<output is skipped here>

LINUXST:/usr/local/src/maildrop-1.8.1 # cd ..
LINUXST:/usr/local/src #

```

Figure 4-11 Installing maildrop

## 6. Installing qmail-scanner-1.25

Now the final step is to install qmail-scanner. Since qmailscanner is to integrate Clam Antivirus and SpamAssassin with qmail, you should install and configure both Clam Antivirus and SpamAssassin before installing qmailscanner.

**Note:** Due to this reason, qmailscanner is not discussed here. You can find information about the installation and configuration of qmail-scanner in 4.4.6, “Installation and configuration of qmail-scanner” on page 224.

## 4.2.4 Configuring qmail

qmail configuration involves creating a few scripts and setting up them. The following section describes the necessary qmail configuration.

### 1. About qmail configuration files

Unlike other packages, qmail uses one file each for every variable with the variable name as the file name and the content in the file as the value for the variable. These files are called control files since they are placed in the qmail configuration folder control located at /var/qmail/control. qmail has many such files. Not all files are created or needed by default. It depends on our requirement. If you want to include or enable a variable then you have to create a file with a variable name as a file name and a value for that variable as the content of the file. By default, only a few such files are created. Figure 4-12 on page 196 shows the qmail configuration files created by default in our qmail installation so far.

```

LINUXST:/var/qmail/control# ls -la
total 20
drwxr-xr-x 2 root qmail 1024 May 11 10:08 .
drwxr-xr-x 10 root root 1024 May 11 02:48 ..
-rw-r--r-- 1 root root 21 May 11 02:40 defaultdomain
-rw-r--r-- 1 root root 26 May 11 02:40 locals
-rw-r--r-- 1 root root 26 May 11 02:40 me
-rw-r--r-- 1 root root 21 May 11 02:40 plusdomain
-rw-r--r-- 1 root root 26 May 11 02:40 rcpthosts

```

Figure 4-12 qmail control files

## 2. Creating qmail aliases

qmail uses files for every alias. You have to create an alias for root. This is because root does not receive mail in qmail. In fact, qmail-lspawn will immediately die and bounce the message if UID == 0. Type the following to create all necessary aliases.

```

cd /var/qmail/alias
touch .qmail-root
touch .qmail-postmaster
touch .qmail-mailer-daemon

```

## 3. Setting up daemontools and creating control directories

Now we have to set up so that qmail starts when the system comes up. This is taken care of by the daemontools package.

- a. First create the qmail service and log directories by entering the commands given below.

```

mkdir -p /var/qmail/supervise/qmail-send/log
mkdir -p /var/qmail/supervise/qmail-smtpd/log
chmod +t /var/qmail/supervise/qmail-send
chmod +t /var/qmail/supervise/qmail-smtpd
mkdir -p /var/log/qmail/qmail-send
mkdir -p /var/log/qmail/qmail-smtpd
chown -R qmail /var/log/qmail

```

- b. Now we are going to create couple of scripts under /var/qmail/supervise/qmail-send and qmail-smtpd.

- i. Create a script called *run* under /var/qmail/supervise/qmail-send by executing the command:

```
vi /var/qmail/supervise/qmail-send/run
```

and include the following lines in the file:

```
#!/bin/sh
exec /var/qmail/rc
```

Save the file and exit.

- ii. Create a similar script under `/var/qmail/supervise/qmail-send/log` by executing the command:

```
vi /var/qmail/supervise/qmail-send/log/run
```

and include the following lines in the file:

```
#!/bin/sh
The following line should be in the same line.
exec /usr/local/bin/setuidgid qmail /usr/local/bin/multilog t s2500000
/var/log/qmail/qmail-send
```

Save the file and exit.

- iii. Now create another script called `run` under `/var/qmail/supervise/qmail-smtpd` by executing the command:

```
vi /var/qmail/supervise/qmail-smtpd/run
```

and include the following lines in the file:

```
#!/bin/sh
#QMAILQUEUE="/var/qmail/bin/qmail-scanner-queue" export QMAILQUEUE
#Note: The above line should be uncommented while using qmail-scanner
QMAILDUID=`id -u qmaild`
NOFILESGID=`id -g qmaild`
#The following line should be in the same line
exec /usr/local/bin/softlimit -m 2000000 /usr/local/bin/tcpserver -H -R -v -p -x
/etc/tcp.smtp.cdb -u $QMAILDUID -g $NOFILESGID 0 smtp /var/qmail/bin/qmail-smtpd
2>&1
```

Save the file and exit.

- iv. Now create a similar file under `qmail-smtpd`'s log file by executing the command:

```
vi /var/qmail/supervise/qmail-smtpd/log/run
```

and include the following lines in the file:

```
#!/bin/sh
#The following line should be in the same line
exec /usr/local/bin/setuidgid qmail /usr/local/bin/multilog t s2500000
/var/log/qmail/qmail-smtpd
```

Save the file and exit.

- v. Now make all run files executable by typing the following commands:

```
chmod 755 /var/qmail/supervise/qmail-send/run
chmod 755 /var/qmail/supervise/qmail-send/log/run
chmod 755 /var/qmail/supervise/qmail-smtpd/run
chmod 755 /var/qmail/supervise/qmail-smtpd/log/run
```

#### 4. Access Control

Before starting up the `qmail smtpd`, you need to do some access control. Note that a default configuration does not allow any relaying, even from localhost.

- a. Create a file `tcp.smtp` under `/etc` by executing the command:

```
vi /etc/tcp.smtp
```

To allow relaying from localhost and your local network, you would have to enter something like this:

```
127.0.0.1:allow,RELAYCLIENT=""
192.168.1.:allow,RELAYCLIENT=""
:allow
```

Save the file and exit.

**Note:** Just replace the IP addresses in the example with your real IP addresses. Note that you can have more than one line in your rules file, to allow relaying from various addresses.

- b. After you have created your rules, you need to activate them. The `tcpserver` works by reading a `cdb` (database) file. You use the `tcprules` program to build the database file from your `/etc/tcp.smtp` file. This is accomplished by the following command:

```
tcprules /etc/tcp.smtp.cdb /etc/tcp.smtp.tmp < /etc/tcp.smtp
```

## 5. Creating `/var/qmail/rc`

`/var/qmail/rc` is one of the most important `qmail` files. It is responsible for starting `qmail`, but more importantly, it tells `qmail` what to do with every e-mail it receives. `rc` contains the default delivery instructions for `qmail`. There are different ways of default mail delivery. The most commonly used format is Mailbox format, and it works with most MUAs. Other formats are the traditional mailbox (`/var/spool/mail`) and Maildir. Since those have limitations in speed and have less MUA support, respectively, we use Mailbox. In this format the mail is delivered in the file `$HOME/.Mailbox`.

- a. Type the following command to create the file `defaultdelivery` under `qmail`'s configuration directory and make the default delivery option:

```
echo ./Mailbox >/var/qmail/control/defaultdelivery
```

- b. Now create `qmail`'s `rc` file by executing the command:

```
vi /var/qmail/rc
```

and add the following lines:

```
#!/bin/sh
Using control/defaultdelivery from qmail-local to deliver messages by default
exec env - PATH="/var/qmail/bin:$PATH" \
qmail-start "`cat /var/qmail/control/defaultdelivery`"
```

Then save the file and exit.

**Note:** This script uses backquotes (```), not single quotes (`'`).

- c. Now make the `rc` script executable by changing the file permissions.

```
chmod 755 /var/qmail/rc
```

## 6. Including the virtual IP in the `rcpthosts` list

Since all the incoming mail is routed via `LINUXFW`, the `rcpthosts` should list the virtual IP of `LINUXFW` to accept mail from. Just add a line that reads the virtual IP of the `LINUXFW` partition in the `rcpghost` file located at `/var/qmail/control`. The steps are described here.

- a. Edit the file `rcpthosts` by typing:

```
vi /var/qmail/control/rcpthosts
```

- b. Add the virtual IP of `LINUXFW` at the end of the file. Refer to the planning worksheet for the IP details. In our setup, we use `10.1.1.10` for the virtual IP of `LINUXFW`. Hence, we have added the line as shown below:

```
10.1.1.10
```

- c. Save the file and quit the editor.

## 7. Configuring for local delivery

- Case 1: If you are using Domino with `qmail`

For Domino users the mail should be delivered to the i5 OS partition, once they have completed all filtering and scanning processes. This requires a change in the file locals, which is a qmail configuration file.

- i. Edit the file locals located at /var/qmail/control by typing:

```
vi /var/qmail/control/locals
```

- ii. Remove all the existing lines.
- iii. Add the virtual IP of the i5 OS partition. In our setup, we use 10.1.1.30 as the virtual IP of the i5 OS. Hence, we have added the line shown below.

```
10.1.1.30
```

- iv. Save the file and quit the editor.

- Case 2: If you are using qmail only

In this setup the mail is to be delivered locally. So the file locals, a qmail configuration file, should read the local IP of the LINUXST partition.

- i. Edit the file locals located at /var/qmail/control by typing:

```
vi /var/qmail/control/locals
```

- ii. Remove all the existing lines.
- iii. Add the virtual IP of LINUXST partition. In our setup, we use 10.1.1.20 as the virtual IP of LINUXST. Hence, we have added the line shown below.

```
10.1.1.20
```

- iv. Save the file and quit the editor.

## 8. Starting qmail

Everything is in place now, and we are ready to start qmail. Since svscan is already running, all we need to do is create symlinks in /service to our control directories in /var/qmail/supervise.

**Note:** The service *svscan* and the directory */service* are created by daemontools.

- a. Create the symlinks by executing the command:

```
ln -s /var/qmail/supervise/* /service/
```

- b. qmail services should start within a few seconds. To check if the services have started, you can run ps, as shown below.

```
ps auxww | grep qmail
```

You should see several tasks running, at the very least qmail-send, and some supervise processes similar to the one shown in Figure 4-13 on page 200.

```

LINUXST:/var/qmail/alias # ln -s /var/qmail/supervise/* /service/
LINUXST:/var/qmail/alias # ps auxww | grep qmail
root 5329 0.0 0.0 1372 348 ? S 14:31 0:00 supervise qmail-smtpd
qmail 5331 0.0 0.0 1384 348 ? S 14:31 0:00 /usr/local/bin/multilog t
s2500000 /var/log/qmail/qmail-smtpd
qmaild 5332 0.0 0.0 1644 532 ? S 14:31 0:00 /usr/local/bin/tcpserver
-H -R -v -p -x /etc/tcp.smtp.cdb -u 1004 -g 1000 0 smtp /var/qmail/bin/qmail-smtpd
root 5333 0.0 0.0 1372 352 ? S 14:31 0:00 supervise qmail-send
qmail 5339 0.0 0.0 1384 344 ? S 14:31 0:00 /usr/local/bin/multilog t
s2500000 /var/log/qmail/qmail-send
root 5408 0.0 0.0 1836 700 pts/1 S+ 14:32 0:00 grep qmail
LINUXST:/var/qmail/alias #

```

Figure 4-13 qmail processes

## 9. Creating qmail control script

If you manually execute the `/var/qmail/rc` script, qmail would be partially started. But we want qmail started up automatically every time the system is booted and we want it shut down cleanly when the system is halted. The result is `qmailctl`, which is a custom qmail control script used to start, stop, and restart qmail like most UNIX services.

One such script is available at <http://www.lifewithqmail.org>. We can simply download and use it.

- a. Change to the qmail's binary directory by executing the command:

```
cd /var/qmail/bin
```

- b. Download `qmailctl` by executing the command:

```
curl http://www.lifewithqmail.org/qmailctl-script-dt70 > qmailctl
```

**Note:** The `curl` command does not have an `-O` option here. This is because we want the destination file to be `qmailctl`. This is accomplished by the redirection.

- c. Change the file permission to make this script executable by typing the command:

```
chmod 755 /var/qmail/bin/qmailctl
```

- d. Link this file to `/usr/sbin` by executing the command:

```
ln -s /var/qmail/bin/qmailctl /usr/sbin
```

Figure 4-14 on page 201 shows the installation of `qmailctl`.

```

LINUXST:~ # cd /var/qmail/bin
LINUXST:/var/qmail/bin # curl http://www.lifewithqmail.org/qmailctl-script-dt70 >
qmailctl
 % Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 3053 100 3053 0 0 3083 0 0:00:00 0:00:00 0:00:00 13296
LINUXST:/var/qmail/bin # chmod 755 /var/qmail/bin/qmailctl
LINUXST:/var/qmail/bin # ln -s /var/qmail/bin/qmailctl /usr/sbin

```

Figure 4-14 Installing qmailctl

- e. Done. Now qmail can be started and stopped by simply executing the commands:

```

qmailctl stop
qmailctl start

```

**Note:** By using qmailctl, you can stop, start, pause, continue, reload, flush, and restart the qmail service. For the complete list type `qmailctl help` at the command prompt.

Figure 4-15 on page 202 shows the different outputs of qmailctl.

```

LINUXST:/var/qmail/bin # qmailctl stop
Stopping qmail...
 qmail-smtpd
 qmail-send
LINUXST:/var/qmail/bin # qmailctl start
Starting qmail
LINUXST:/var/qmail/bin # qmailctl help
 stop -- stops mail service (smtp connections refused, nothing goes out)
 start -- starts mail service (smtp connection accepted, mail can go out)
 pause -- temporarily stops mail service (connections accepted, nothing leaves)
 cont -- continues paused mail service
 stat -- displays status of mail service
 cdb -- rebuild the tcpserver cdb file for smtp
restart -- stops and restarts smtp, sends qmail-send a TERM & restarts it
doqueue -- schedules queued messages for immediate delivery
reload -- sends qmail-send HUP, rereading locals and virtualdomains
queue -- shows status of queue
 alrm -- same as doqueue
 flush -- same as doqueue
 hup -- same as reload
LINUXST:/var/qmail/bin # qmailctl stat
/service/qmail-send: up (pid 6501) 313 seconds
/service/qmail-send/log: up (pid 6183) 313 seconds
/service/qmail-smtpd: up (pid 6186) 313 seconds
/service/qmail-smtpd/log: up (pid 6187) 313 seconds
messages in queue: 0
messages in queue but not yet preprocessed: 0
LINUXST:/var/qmail/bin #

```

Figure 4-15 Working with qmailctl

## 10. Testing the installation

You can run `qmailctl` to verify that the services are up and running. At the command prompt type:

```
qmailctl stat
```

You might see an output similar to the following.

```

/service/qmail-send: up (pid 6501) 313 seconds
/service/qmail-send/log: up (pid 6183) 313 seconds
/service/qmail-smtpd: up (pid 6186) 313 seconds
/service/qmail-smtpd/log: up (pid 6187) 313 seconds
messages in queue: 0
messages in queue but not yet preprocessed: 0

```

If all four services are *up* for more than a second, qmail is running fine.

## 4.3 Clam AntiVirus

Clam AntiVirus is a free, open source virus scanner that uses a frequently updated database of virus definitions. Clam AntiVirus scans local files, as well as all incoming e-mail messages. It detects over 30,000 viruses, worms, and Trojan horses, including Microsoft® Office and MacOffice macro viruses. It scans within archives and compressed files (and also protects against archive bombs). Built-in support includes:

- ▶ Zip files
- ▶ RAR (2.0) files

- ▶ Tar files
- ▶ Gzip files
- ▶ Bzip2 files
- ▶ MS OLE2 files
- ▶ MS Cabinet files
- ▶ MS CHM (Compiled HTML) files
- ▶ MS SZDD compression format files
- ▶ UPX (all versions) files
- ▶ FSG (1.3, 1.31, 1.33, 2.0) files
- ▶ Petite (2.x) files

### 4.3.1 Installing Clam Antivirus

Clam AntiVirus uses the compression techniques contained within the zLibraries, so the zlib package must be installed before Clam AntiVirus can be used. To enable Clam AntiVirus to scan mail, ClamSMTP must be installed to transport mail between MTAs, such as Postfix and gmail, and Clam AntiVirus. Clam AntiVirus and zlib are available built-in with SLES9. If you have selected those packages to be installed while installing the SLES9, they must be upgraded to the latest version for better results. In case of RHEL4, Clam AntiVirus is not included in the minimal installation; only zlib is included. But the installed version of zlib must be upgraded to the latest version.

#### Installing zlib 1.2.2

The zlib is a general purpose compression library that provides in-memory compression and decompression functions, including integrity checks of the uncompressed data. The library also supports reading and writing files in gzip (.gz) format with an interface.

1. The first step in the zlib installation process is to download the source. With the virtual console still open, make sure you are in the /usr/local/src/ directory by executing the command:

```
cd /usr/local/src
```

2. Create a sectools directory by executing the command:

```
mkdir sectools
```

3. Change to the sectools directory by executing the command:

```
cd sectools
```

4. To download the source type at the command line:

```
curl -O http://www.zlib.net/zlib-1.2.2.tar.gz
```

Or:

```
curl -O http://www.gzip.org/zlib/zlib-1.2.2.tar.gz
```

5. Now the downloaded source must be extracted. At the command line type:

```
tar xzf zlib-1.2.2.tar.gz
```

6. Change to the zlib directory by executing the command:

```
cd zlib-1.2.2
```

7. Now the package is configured by executing the command:

```
./configure
```

8. Make the executables and place them in the directories they are meant for by executing the commands:

```
make; make install
```

9. Exit the directory by executing the command:

```
cd ..
```

Figure 4-16 shows the steps involved in installing zlib-1.2.2.

```
LINUXST:~ # cd /usr/local/src
LINUXST:/usr/local/src # mkdir sectools
LINUXST:/usr/local/src # cd sectools
LINUXST:/usr/local/src/sectools # curl -O http://www.zlib.net/zlib-1.2.2.tar.gz
 % Total % Received % Xferd Average Speed Time Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 8408 0 0:00:51 0:00:51 0:00:00 0k
LINUXST:/usr/local/src/sectools # tar xzf zlib-1.2.2.tar.gz
LINUXST:/usr/local/src/sectools # cd zlib-1.2.2/
LINUXST:/usr/local/src/sectools/zlib-1.2.2 # ./configure
Checking for gcc...
Building static library libz.a version 1.2.2 with gcc.
Checking for unistd.h... Yes.
Checking whether to use vs[n]printf() or s[n]printf()... using vs[n]printf()
Checking for vsnprintf() in stdio.h... Yes.
Checking for return value of vsnprintf()... Yes.
Checking for errno.h... Yes.
Checking for mmap support... Yes.
LINUXST:/usr/local/src/sectools/zlib-1.2.2 # make; make install
gcc -O3 -DUSE_MMAP -c -o example.o example.c
gcc -O3 -DUSE_MMAP -c -o adler32.o adler32.c

<the output is skipped here>

cp zlib.3 /usr/local/share/man/man3
chmod 644 /usr/local/share/man/man3/zlib.3
LINUXST:/usr/local/src/sectools/zlib-1.2.2 # cd ..
LINUXST:/usr/local/src/sectools #
```

Figure 4-16 Installing zlib

### Installing GMP-4.1.4

The GNU MP is a library for arbitrary precision arithmetic, operating on signed integers, rational numbers, and floating point numbers. It has a rich set of functions, and the functions have a regular interface.

**Important:** The package GMP (GNU MP) is optional but highly recommended because it allows freshclam to verify the digital signatures of the virus databases.

If freshclam was compiled without GMP support it will display "SECURITY WARNING: NO SUPPORT FOR DIGITAL SIGNATURES" on every update.

You can download and install GNU MP following the quick steps below.

1. Change to the directory by executing the command:

```
cd /usr/local/src/sectools
```

2. Download the GMP source by executing the command:
 

```
curl -O ftp://ftp.gnu.org/gnu/gmp/gmp-4.1.4.tar.gz
```
3. Extract the source by executing the command:
 

```
tar xzf gmp-4.1.4.tar.gz
```
4. Change to the directory by executing the command:
 

```
cd gmp-4.1.4
```
5. Install GMP by executing the commands one-by-one, as below.
 

```
./configure
make
make install
```
6. Exit the directory by executing the command:
 

```
cd ..
```

Figure 4-17 shows the steps involved in installing GMP-4.1.4.

```
LINUXST:~ # cd /usr/local/src
LINUXST:/usr/local/src # mkdir sectools
LINUXST:/usr/local/src # cd sectools
LINUXST:/usr/local/src/sectools # curl -O ftp://ftp.gnu.org/gnu/gmp/gmp-4.1.4.tar.gz
% Total % Received % Xferd Average Speed Time
Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 0k 0 0:00:03 0:00:03 0:00:00 0k
LINUXST:/usr/local/src/sectools # tar xzf zlib-1.2.2.tar.gz
LINUXST:/usr/local/src/sectools # cd gmp-4.1.4/
LINUXST:/usr/local/src/sectools/gmp-4.1.4 # ./configure
checking build system type... powerpc64-unknown-linux-gnu
checking host system type... powerpc64-unknown-linux-gnu

<the output is skipped here>

LINUXST:/usr/local/src/sectools/gmp-4.1.4 # make

<output is skipped>output is skipped here>

LINUXST:/usr/local/src/sectools/gmp-4.1.4 # make install

<output is skipped here>

LINUXST:/usr/local/src/sectools/gmp-4.1.4 # cd ..
LINUXST:/usr/local/src/sectools #
```

Figure 4-17 Installing GMP

## Installing Clam AntiVirus 0.85

This is the latest version released at the time of writing this redpaper. This version of ClamAV improves detection of JPEG-based exploits and introduces support for TNEF files and new detection mechanisms. Various bugfixes (including problems with scanning digest mail files) and improvements have been made.

1. As the zLibrary is installed, installation of the Clam AntiVirus system can begin. First, add a Clam AntiVirus user and group to the system. With the virtual console still open, at the command line type:
 

```
groupadd clamav
```

2. To add the user in the group clamav, at the command line type:
 

```
useradd -g clamav -s /bin/false -c "Clam AntiVirus" clamav
```
  3. Change to the directory sectools by executing the command:
 

```
cd /usr/local/src/sectools
```
  4. Next, download the Clam AntiVirus 0.85clamav-0.85 source code. At the command line type:
 

```
curl -O http://easynews.dl.sourceforge.net/sourceforge/clamav/clamav-0.85.tar.gz
```
  5. Now the source must be extracted. At the command line type:
 

```
tar xzf clamav-0.85.tar.gz
```
  6. The extracted source must now be compiled and installed. At the command line type the following, one-by-one.
 

```
cd clamav-0.85
./configure --sysconfdir=/etc
make
make install
```
  7. Now copy the clamd executable to /etc/init.d so that clamd can be started/stopped as other daemons. Type the following commands:
 

```
cd clamd
cp clamd /etc/init.d
```
  8. Exit the directory by executing the command:
 

```
cd ../../
```
- Figure 4-18 shows the steps we used while installing.

```
LINUXST:~ # groupadd clamav
LINUXST:~ # useradd -g clamav -s /bin/false -c "Clam AntiVirus" clamav
LINUXST:~ # cd /usr/local/src/sectools
LINUXST:/usr/local/src/sectools # curl -O
http://easynews.dl.sourceforge.net/sourceforge/clamav/clamav-0.85.tar.gz
 % Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 0k 0 0:00:11 0:00:11 0:00:00 0k
LINUXST:/usr/local/src/sectools # tar xzf clamav-0.85.tar.gz
LINUXST:/usr/local/src/sectools # cd clamav-0.85/
LINUXST:/usr/local/src/sectools/clamav-0.85 # ./configure --sysconfdir=/etc

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamav-0.85 # make

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamav-0.85 # make install

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamav-0.85 # cd clamd
LINUXST:/usr/local/src/sectools/clamav-0.85/clamd # cp clamd /etc/init.d
LINUXST:/usr/local/src/sectools/clamav-0.85/clamd # cd ../../
LINUXST:/usr/local/src/sectools #
```

Figure 4-18 Installing Clam Antivirus

## Installing ClamSMTP 1.4.1

The final step in installing the Clam AntiVirus system is to install ClamSMTP. ClamSMTP is a filter that works in conjunction with Clam AntiVirus. ClamSMTP intercepts all incoming mail and passes the messages to Clam AntiVirus so that they may be scanned before being passed back to the mail server.

1. With the virtual console still open, change to the present working directory by executing the command:

```
cd /usr/local/src/sectools
```

2. Now download the ClamSMTP 1.4.1 tarball. At the command line type:

```
curl -O http://memberwebs.com/nielsen/software/clamsmtp/clamsmtp-1.4.1.tar.gz
```

3. Now extract the downloaded source. At the command line type:

```
tar xzf clamsmtp-1.4.1.tar.gz
```

4. Change to the ClamSMTP directory by executing the command:

```
cd clamsmtp-1.4.1
```

5. Now compile the source and install ClamSMTP. At the command line type the following, one by one:

```
./configure
```

```
make
```

```
make install
```

6. Copy the configuration file by executing the command:

```
cp /usr/local/src/sectools/clamsmtp-1.4.1/doc/clamsmtpd.conf /etc/
```

7. Exit the directory by executing the command:

```
cd ..
```

The Clam AntiVirus system is now installed with all the components that it requires to protect your system and your e-mail from viruses. Figure 4-19 on page 208 shows the steps involved in installing ClamSMTP.

```

LINUXST:/var/qmail/bin # cd /usr/local/src/sectools
LINUXST:/usr/local/src/sectools # curl -O
http://memberwebs.com/nielsen/software/clamsmtp/clamsmtp-1.4.1.tar.gz
 % Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 12283 0 0:00:11 0:00:11 0:00:00 0k
LINUXST:/usr/local/src/sectools # tar zxf clamsmtp-1.4.1.tar.gz
LINUXST:/usr/local/src/sectools # cd clamsmtp-1.4.1/
LINUXST:/usr/local/src/sectools/clamsmtp-1.4.1 # ./configure

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamsmtp-1.4.1 # make

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamsmtp-1.4.1 # make install

<output is skipped here>

LINUXST:/usr/local/src/sectools/clamsmtp-1.4.1 # cd ..
LINUXST:/usr/local/src/sectools #

```

Figure 4-19 Installing ClamSMTP

**Note:** The procedure to install and configure Clam AntiVirus on RHEL4 is the same as for SLES9.

### 4.3.2 Configuring Clam AntiVirus for Postfix

Clam AntiVirus uses one config file (located at `/etc/clamd.conf`) to define its behavior. There are many options that may be changed to alter the way Clam AntiVirus functions. For the most part, however, the default options are fine for protecting your system.

1. Edit the Clam AntiVirus config file.
  - a. With the virtual console still open, at the command line type:

```
vi /etc/clamd.conf
```

- b. Comment out the line that reads:

Example

so that it reads:

```
Example
```

**Important:** If the option Example is not commented, clamd will not run.

- c. Change the line that reads:

```
LocalSocket /tmp/clamd
```

to:

```
LocalSocket /etc/clamd.sock
```

- d. Uncomment the lines:

```
FixStaleSocket
```

```
LogFile
LogTime
LogSyslog
User clamav
```

- e. If you want to include your internal mails (that is, mails from local users) to be scanned, then uncomment the following line (recommended):

```
ScanMail
```

- f. Save the file and quit the editor.

2. Next, start Clam AntiVirus. At the command line type:

```
/etc/init.d/clamd start
```

## Configuring FreshClam

FreshClam is a function that is built into Clam AntiVirus. FreshClam updates the virus definitions that Clam AntiVirus uses via the Internet.

1. Create a log file to log all Clam AntiVirus updates.

- a. With the virtual console still open, at the command line type:

```
touch /var/log/clam-update.log
```

- b. Make this log file so that it can be read and written by the user clamav. At the command line type:

```
chmod 600 /var/log/clam-update.log
```

- c. Change the file ownership to clamav. At the command line type:

```
chown clamav:clamav /var/log/clam-update.log
```

2. Edit the configuration file of freshclam, which is located at /etc/freshclam.conf.

- a. At the command line type:

```
vi /etc/freshclam.conf
```

- b. Comment out the line that reads:

```
Example
```

```
so that it reads:
```

```
Example
```

Save the file and quit the editor.

3. Restart the clamd so that the configuration changes will be in effect.

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

4. Now, run the FreshClam daemon to update the databases with the latest virus definitions. At the command line type:

```
freshclam -l /var/log/clam-update.log
```

which may give the output similar that below.

```
ClamAV update process started at Thu May 12 15:58:30 2005
main.cvd is up to date (version: 31, sigs: 33079, f-level: 4, builder: tkojm)
daily.cvd updated (version: 875, sigs: 1218, f-level: 4, builder: diego)
```

This means your mail server is able to get the updates.

5. As you think now, this update should be done in regular intervals, like everyday. By setting up a cron entry in the crontab, this can be accomplished.

- a. As root user, edit your crontab by typing the command:
 

```
crontab -e
```
- b. Add a line as given below, which executes the freshclam to run quietly everyday by 2:30 a.m. for the latest updates on virus definitions.
 

```
30 2 * * * /usr/bin/freshclam --quiet -l /var/log/clam-update.log
```
- c. Save the file and exit.

## Configuring ClamSMTP

ClamSMTP needs other configuration files to point to it apart from extra configuration of it is own. ClamSMTP must be set to run at system startup.

1. Edit the ClamSMTP configuration file.
  - a. The configuration file is located at `/etc/clamsmtpd.conf`. At the command line type:
 

```
vi /etc/clamsmtpd.conf
```
  - b. Modify the line that reads:
 

```
OutAddress: 10026
```

to:

```
OutAddress: 127.0.0.1:10026
```
  - c. Uncomment the line that reads:
 

```
User: clamav
```
  - d. Save the file and quit the editor.
2. Edit the Postfix config files.
  - a. First edit the file `main.cf`. With the virtual console still open, at the command line type:
 

```
vi /etc/postfix/main.cf
```
  - b. Modify the file that reads:
 

```
content_filter =
```

to:

```
content_filter = scan:127.0.0.1:10025
```

The `content_filter` tells Postfix to send all mail through the service called `scan` on port 10025.
  - c. Add a line that reads:
 

```
receive_override_options = no_address_mappings
```
  - d. Save the file and quit the editor.
  - e. Now edit the file `master.cf`. At a command line type:
 

```
vi /etc/postfix/master.cf
```
  - f. Add the following lines at the end of the file. Make sure there are no spaces around the equal signs or commas in the text you added. The first two lines create the scan service. The rest set up a service for accepting mail back into Postfix. All the options prevent a mail loop as well as relax address checking and the like, since that has already been done.
 

```
AV scan filter (used by content_filter)
scan unix - - n - 5 smtp
-o smtp_send_xforward_command=yes
For injecting mail back into postfix from the filter
127.0.0.1:10026 inet n - n - 5 smtpd
```

```
-o content_filter=
-o receive_override_options=no_unknown_recipient_checks,no_header_body_checks
-o smtpd_helo_restrictions=
-o smtpd_client_restrictions=
-o smtpd_sender_restrictions=
-o smtpd_recipient_restrictions=permit_mynetworks,reject
-o mynetworks_style=host
-o smtpd_authorized_xforward_hosts=127.0.0.0/8
```

g. Save the file and quit the editor.

3. Now start up the clamsmtpd by executing the command:

```
clamsmtpd -f /etc/clamsmtpd.conf -p /etc/clamd.sock
```

### 4.3.3 Configuring Clam Antivirus for gmail

Clam AntiVirus uses one config file (located at `/etc/clamd.conf`) to define its behavior. There are many options that may be changed to alter the way that Clam AntiVirus functions. For the most part, however, the default options are fine for protecting your system.

1. Edit the Clam AntiVirus config file.

a. With the virtual console still open, at the command line type:

```
vi /etc/clamd.conf
```

b. Comment out the line that reads:

```
Example
```

so that it reads:

```
Example
```

**Important:** If the option `Example` is not commented, `clamd` will not run.

c. Change the line that reads:

```
LocalSocket /tmp/clamd
```

to:

```
LocalSocket /etc/clamd.sock
```

d. Uncomment the lines:

```
FixStaleSocket
```

```
LogFile
```

```
LogTime
```

```
LogSyslog
```

```
User clamav
```

e. If you want to include your internal mails (that is, mails from local users) to be scanned, then uncomment the following line (recommended):

```
ScanMail
```

f. Save the file and quit the editor.

2. Next, start Clam AntiVirus. At the command line type:

```
/etc/init.d/clamd start
```

### Configuring FreshClam

FreshClam is a function that is built into Clam AntiVirus. FreshClam updates the virus definitions that Clam AntiVirus uses via the Internet.

1. Create a log file to log all Clam AntiVirus updates.
  - a. With the virtual console still open, at the command line type:
 

```
touch /var/log/clam-update.log
```
  - b. Make this log file so that it can be read and written by the user clamav. At the command line type:
 

```
chmod 600 /var/log/clam-update.log
```
  - c. Change the file ownership to clamav. At the command line type:
 

```
chown clamav:clamav /var/log/clam-update.log
```
2. Edit the configuration file of freshclam, which is located at /etc/freshclam.conf.
  - a. At the command line type:
 

```
vi /etc/freshclam.conf
```
  - b. Comment out the line that reads:
 

```
Example
```

 so that it reads:
 

```
Example
```

 Save the file and quit the editor.
3. Restart the clamd so that the configuration changes will be in effect.
 

```
/etc/init.d/clamd stop
/etc/init.d/clamd start
```
4. Now, run the FreshClam daemon to update the databases with the latest virus definitions. At the command line type:
 

```
freshclam -l /var/log/clam-update.log
```

 which may give the output similar to below.
 

```
ClamAV update process started at Thu May 12 15:58:30 2005
main.cvd is up to date (version: 31, sigs: 33079, f-level: 4, builder: tkojm)
daily.cvd updated (version: 875, sigs: 1218, f-level: 4, builder: diego)
```

 This means your mail server is able to get the updates.
5. As you know now, this update should be done in regular intervals, like everyday. By setting up a cron entry in the crontab, this can be accomplished.
  - a. As root user, edit your crontab by typing the command:
 

```
crontab -e
```
  - b. Add a line as given below, which executes the freshclam to run quietly everyday by 2:30 a.m. for the latest updates on virus definitions.
 

```
30 2 * * * /usr/bin/freshclam --quiet -l /var/log/clam-update.log
```
  - c. Save the file and exit.
6. Test clamd.
 

Clam AntiVirus can be tested by scanning a directory. Figure 4-20 on page 213 shows the typical test result of clamd scan.

```

LINUXST:/usr/local/src/sectools # clamscan clamav-0.85
/usr/local/src/sectools/clamav-0.85/test/clam.cab: ClamAV-Test-File FOUND
/usr/local/src/sectools/clamav-0.85/test/clam.exe: ClamAV-Test-File FOUND
/usr/local/src/sectools/clamav-0.85/test/clam.zip: ClamAV-Test-File FOUND
/usr/local/src/sectools/clamav-0.85/test/clam.exe.bz2: ClamAV-Test-File FOUND

----- SCAN SUMMARY -----
Infected files: 4
Time: 8.394 sec (0 m 8 s)
LINUXST:/usr/local/src/sectools #

```

Figure 4-20 Testing clamd

### 4.3.4 Adding Clam to system boot

Clam AntiVirus, FreshClam, and ClamSMTP all need to be started when the system boots up.

1. For SLES9 users, with the virtual console still open, edit the boot.local file. At the command line type:

```
vi /etc/rc.d/boot.local
```

2. Add the following lines to the end of the file.

```

/usr/local/sbin/clamd
/usr/local/bin/freshclam -d -c 2 -l /var/log/clam-update.log
/usr/local/sbin/clamsmtpd -f /etc/clamsmtpd.conf -p /etc/clamd.sock

```

3. Save the file and quit the editor.

**Note:** In the case of RHEL4, edit the file /etc/rc.d/rc.local instead of /etc/rc.d/boot.local.

Clam AntiVirus is now installed and running, and will start at system boot. FreshClam will attempt to update your virus definitions twice a day. All update attempts are written to a log file located at /var/log/clam-update.log.

## 4.4 SpamAssassin

SpamAssassin is an e-mail filter program released freely under the Apache License. It filters spams based on content-matching rules, which also supports DNS-based, checksum-based, and statistical filtering, supported by external programs and online databases. In addition, SpamAssassin has a modular architecture that allows other technologies to be quickly wielded against spam and is designed for easy integration into virtually any e-mail system.

SpamAssassin is generally regarded as one of the most effective spam filters, especially when used in combination with spam databases. Even simple text-matching alone may, for most users, be sufficient to correctly classify a majority of incoming mail.

## Overview of spam

Spam is any message or posting, regardless of its content, that is sent to multiple recipients who have not specifically requested the message. Spam can also be multiple postings of the same message to newsgroups or list servers that are not related to the topic of discussion. Other common terms for spam include unsolicited commercial e-mail, unsolicited bulk e-mail, and junk mail.

E-mail addresses may be harvested from a number of sources. A popular method has been to use e-mail addresses that their owners have published for other purposes. Usenet posts, especially those in archives such as Mail Groups, are a frequent target. Simply searching the Web for pages with addresses—such as corporate staff directories—can yield thousands of addresses, most of them deliverable. Spammers have also subscribed to discussion mailing lists for the purpose of gathering the addresses of posters. The DNS and WHOIS systems require the publication of technical contact information for all Internet domains; spammers have illegally trawled these resources for e-mail addresses. Many spammers utilize programs called Web spiders to find e-mail addresses on Web pages. Also, Web sites that are less reputable or have privacy policies that are less strict may sell the e-mail addresses to spammers. Those who send spam typically have collected a list of e-mail addresses from these many ways. They send messages from numerous different (and often falsified) addresses to all areas of the Web, which makes them hard to track. A single spam run may target tens of millions of possible addresses—many of which are invalid, malformed, or undeliverable.

**Note:** E-mail that is not spam is often called *ham*.

### 4.4.1 Installing SpamAssassin

Since SpamAssassin is a perl module it is installed using CPAN. Alternatively, it can be installed by downloading the source code, and extracting and configuring as with other packages. We prefer to install using CPAN, as it is not only easy to install but also very effective if it requires any other required perl modules. We discuss installation of SpamAssassin in three parts:

- ▶ Installing required Perl modules
- ▶ Installing Mail::SpamAssassin Perl module
- ▶ Installing optional but recommended programs

These are expanded upon below:

- ▶ Installing required Perl modules

We are going to install the required Perl modules for SpamAssassin via CPAN.

- a. As a root user, at the command prompt type:

```
perl -MCPAN -e shell
```

**Note:** If this is the very first time you are accessing CPAN from your server then you have to configure the CPAN module, but perl itself would do autoconfiguration. This is recommended, too. In this case, your screen will end up with a question and waiting for reply like:

```
Are you ready for manual configuration? [yes]
```

```
Let Perl do the auto configuration by typing no here.
```

Now the screen looks like similar to Figure 4-21. Also, the command prompt becomes:

```
cpan>
```

This means you are in the CPAN shell.

```
LINUXST:~ # perl -MCPAN -e shell
Terminal does not support AddHistory.

cpan shell -- CPAN exploration and modules installation (v1.7601)
ReadLine support available (try 'install Bundle::CPAN')

cpan>
```

Figure 4-21 Logging into CPAN shell

- b. It is often the case that a module you want to install may require other modules for its proper operation, and perhaps one or more of these additional modules have not yet been installed. The CPAN system includes a way to check to see what other modules are required, and you can configure it to automatically follow and install missing prerequisites. Especially with a large collection of modules like SpamAssassin, you may expect to see such prerequisites crop up. To handle the prerequisites, automatically configure the CPAN session to follow. At the cpan prompt type:

```
cpan> o conf prerequisites_policy follow
```

**Note:** This `cpan` command overrides the default configuration. Note that `o` means override.

- c. Installing HTML::Parser.

As its name says, this Perl module for recognizing HTML markup languages. As HTML is used for an ever-increasing amount of e-mail, this dependency is unavoidable.

To install HTML::Parser, at the cpan prompt type:

```
cpan> install HTML::Parser
```

- d. Installing Net::DNS.

This Perl module is used for all DNS-based tests (SBL, XBL, SpamCop, DSBL, etc.), to perform MX checks, and is also used when manually reporting spam to SpamCop.

To install Net::DNS, at the cpan prompt type:

```
cpan> install Net::DNS
```

It will ask where do you want to do some tests. The default answer is yes. Just press the Enter key to proceed.

e. Installing Mail::SPF::Query.

This Perl module is used to check DNS Sender Policy Framework (SPF) records to fight e-mail address forgery and make it easier to identify spam.

To install Mail::SPF::Query, at the cpan prompt type:

```
cpan> install Mail::SPF::Query
```

f. Installing IP::Country::Fast.

This Perl module is used by the RelayCountry plugin (not enabled by default) to determine the domain country codes of each relay in the path of an e-mail.

To install IP::Country::Fast, at the cpan prompt type:

```
cpan> install IP::Country::Fast
```

g. Installing Digest::Nilsimsa.

This module is the Perl version of Nilsimsa code, actually required for Razor, one of the optional program we are going to install later.

To install Digest::Nilsimsa, at the cpan prompt type:

```
cpan> install Digest::Nilsimsa
```

Figure 4-22 shows the installation of Perl modules.

```
cpan> o conf prerequisites_policy follow
prerequisites_policy follow

cpan> install HTML::Parser

<output is skipped here>

cpan> install Net::DNS

<output is skipped here>

cpan> install Mail::SPF::Query

<output is skipped here>

cpan> install IP::Country::Fast

<output is skipped here>

cpan> install Digest::Nilsimsa

<output is skipped here>

cpan>
```

Figure 4-22 Installing required Perl modules

We have installed all required Perl modules. Now proceed to install SpamAssassin.

► Installing Mail::SpamAssassin Perl module

Like other Perl modules, Mail::SpamAssassin can be installed by typing at the cpan prompt:

```
cpan> install Mail::SpamAssassin
```

The screen now looks similar to Figure 4-23.

```
cpan> install Mail::SpamAssassin
CPAN: Storable loaded ok
Going to read /root/.cpan/Metadata
Warning: Found only 0 objects in /root/.cpan/Metadata
LWP not available
CPAN: Net::FTP loaded ok
Fetching with Net::FTP:
 ftp://ftp.perl.org/pub/CPAN/authors/01mailrc.txt.gz

<remaining output is skipped here>

cpan> quit
Terminal does not support GetHistory.
Lockfile removed.
LINUXST:~ #
```

Figure 4-23 Installing SpamAssassin using CPAN

- ▶ Since we have installed all required Perl modules we can exit from the cpan prompt by typing:

```
cpan> quit
```

Now you are back at the OS command prompt.

- ▶ Installing optional but recommended programs

There are some non-Perl programs that can increase the overall effectiveness of SpamAssassin if they are installed. We recommend installing those programs. The non-Perl programs we are going to install are:

- Razor
- DCC (Distributed Checksum Clearinghouse)
- Pyzor

To install those:

- Installing Razor

Razor is a distributed, collaborative, spam detection and filtering network. Razor establishes a distributed and constantly updating catalogue of spam in propagation that is consulted by e-mail clients to filter out known spam.

Razor has two different packages, namely, Razor-agents and Razor-agents-sdk. The SDK packages bundle with all required Perl modules for Razor. Since we have already installed the required modules we can directly install the main package, Razor-agents. The following steps describe the installation of Razor-agents.

- i. Change to the directory /usr/local/src/sectools by typing:

```
cd /usr/local/src/sectools
```

ii. Download Razor-agents from sourceforge.net by typing:

```
curl -O
http://easynews.dl.sourceforge.net/sourceforge/razor/razor-agents-2.67.tar.gz
```

iii. Extract the tar file by typing:

```
tar xzf razor-agents-2.67.tar.gz
```

iv. Change the directory to razor-agents-2.67 by typing:

```
cd razor-agents-2.67
```

v. Now execute the following commands to complete the installation.

```
perl Makefile.PL
make
make install
```

vi. If you see no errors, Razor is installed. Exit the directory by typing:

```
cd ..
```

#### – Installing DCC

Distributed Checksum Clearinghouse (DCC) is a system similar to Razor. It supports fuzzy checksums and therefore detects more spam than Razor does at the moment.

The following steps describe the installation of DCC.

i. Change to the directory /usr/local/src/sectools by typing:

```
cd /usr/local/src/sectools
```

ii. Download the source from <http://www.rhyolite.com/anti-spam/dcc> by executing the command:

```
curl -O http://www.dcc-servers.net/dcc/source/dcc-dccproc.tar.Z
```

iii. Extract the source from the tar ball by typing:

```
tar xzf dcc-dccproc.tar.Z
```

iv. Get into the dcc source directory by typing:

```
cd dcc-dccproc-*
```

v. Now execute the following commands to complete the installation.

```
./configure
make
make install
```

vi. If it shows no error, DCC is successfully installed. Exit the directory by typing:

```
cd ..
```

#### – Installing Pyzor

Pyzor is a Python implementation of a spam-blocking network system. It is used to check message signatures against the Pyzor collaborative filtering network.

The following steps describe the installation of Pyzor.

i. Change to the directory /usr/local/src/sectools by typing:

```
cd /usr/local/src/sectools
```

ii. Download the source code from sourceforge.net by typing:

```
curl -O http://ufpr.dl.sourceforge.net/sourceforge/pyzor/pyzor-0.4.0.tar.bz2
```

iii. Extract the source tarball by typing:

```
tar xvj pyzor-*.tar.bz2
```

iv. Get into the pyzor source directory by typing:

```
cd pyzor-*
```

v. Now execute the following commands to complete the installation:

```
python setup.py build
python setup.py install
```

vi. If you see no errors, Pyzor is successfully installed. Exit the directory by typing:

```
cd ..
```

vii. Now change the common Pyzor directory permission to world writable by typing:

```
chmod -R a+rX /usr/share/doc/pyzor /usr/lib/python2.2/site-packages/pyzor
```

viii. Now change the user Pyzor binaries to be executable by all users. At the command prompt type:

```
chmod a+rX /usr/bin/pyzor /usr/bin/pyzord
```

## 4.4.2 Overview: Configuration of SpamAssassin

The configuration file of SpamAssassin, `/etc/mail/spamassassin/local.cf`, deals with different options. Each option has its own variables and values. Among all, only few of them are discussed here based on their usage in our configuration file.

### Whitelist and blacklist option

These options are used to either include or exclude the mail addresses from spam scanning.

► `whitelist_from address@domain.com`

Used to specify addresses which send mail that is often tagged (incorrectly) as spam.

Whitelist and blacklist addresses are now file-glob-style patterns. It means `friend@somewhere.com`, `*@isp.com`, or `*.domain.net` will all work. Specifically, `*` and `?` are allowed, but all other metacharacters are not. Regular expressions are not used for security reasons.

Multiple addresses per line, separated by spaces, is OK. Multiple `whitelist_from` lines is also OK. For example:

```
whitelist_from rob@something.com sen@something.com
whitelist_from *@something.com
```

► `unwhitelist_from address@domain.com`

Used to override a default `whitelist_from` entry, so for example a distribution `whitelist_from` can be overridden in a `local.cf` file. It is useful if you include all address from a particular domain (`*@something.com`) but want to exclude few of them.

Example: `unwhitelist_from alex@something.com john@something.com`

**Note:** The `unwhitelist_from` option is useful when you want to include your internal mails in the scanning process. SpamAssassin is not configured to scan the internal mails, mails from the same host, by default. If you want to include the internal mails just use this option like:

```
unwhitelist_from *@localhost *@secureyouremails.com
```

► `blacklist_from address@domain.com`

Used to specify addresses which send mail that is often tagged (incorrectly) as non-spam, but which the user does not want. Same format as `whitelist_from`.

- ▶ `unblacklist_from address@domain.com`

Used to override a default `blacklist_from` entry. Same format as `unwhitelist_from`.

## Scoring options

Every mail gets a value called score while it is being scanned. Based on the score the mail is either considered as spam or not. There are few scoring options, among them the option `required_hits` is important and discussed here.

- ▶ `required_hits n.nn` (default: 5)

Set the number of hits required before a mail is considered spam. `n.nn` can be an integer or a real number. 5.0 is the default setting, and is quite aggressive; it would be suitable for a single-user setup, but if you're an ISP installing SpamAssassin, you should probably set the default to be more conservative, like 8.0 or 10.0. It is not recommended to automatically delete or discard messages marked as spam, as your users will complain, but if you choose to do so, only delete messages with an exceptionally high score such as 15.0 or higher.

## Message tagging options

The options deal with how to display the mail if it is spam.

- ▶ `rewrite_subject { 0 | 1 }` (default: 0)

By default, the subject lines of suspected spam will not be tagged. This can be enabled here.

- ▶ `subject_tag STRING ...` (default: `*****SPAM*****`)

Text added to the Subject: line of mails that are considered spam, if `rewrite_subject` is 1.

- ▶ `report_safe { 0 | 1 | 2 }` (default: 1)

If this option is set to 1, if an incoming message is tagged as spam, instead of modifying the original message, SpamAssassin will create a new report message and attach the original message as a `message/rfc822` MIME part (ensuring the original message is completely preserved, not easily opened, and easier to recover).

If this option is set to 2, then original messages will be attached with a content type of `text/plain` instead of `message/rfc822`. This setting may be required for safety reasons on certain broken mail clients that automatically load attachments without any action by the user. This setting may also make it somewhat more difficult to extract or view the original message.

If this option is set to 0, incoming spam is only modified by adding some X-Spam- headers and no changes will be made to the body. In addition, a header named X-Spam-Report will be added to spam. You can use the `remove_header` option to remove that header after setting `report_safe` to 0.

- ▶ `use_bayes { 0 | 1 }` (default: 1)

Whether to use the naive-Bayesian-style classifier built into SpamAssassin.

## Learning options

This option is used to set the analysis system to determine whether each message is spam based on previous examples of spam and non-spam.

- ▶ `bayes_auto_learn { 0 | 1 }` (default: 1)

Whether SpamAssassin should automatically feed high-scoring mails (or low-scoring mails, for non-spam) into its learning systems. The only learning system supported currently is a naive-Bayesian-style classifier.

## Network test options

These options are set to deal network based testing like checking the blacklist IPs and including filtering networks.

- ▶ `skip_rbl_checks { 0 | 1 }` (default: 0)

By default, SpamAssassin will run RBL checks. If your ISP already does this for you, set this to 1.

**Note:** RBL, Realtime Blackhole List, is a list of IP addresses whose owners refuse to stop the proliferation of spam.

The RBL list is updated in every new version of SpamAssassin. Users do not have an easy option to add a new realtime block IP in the installed one. The file `/usr/share/spamassassin/20_dnsbl_tests.cf` contains configuration for the RBL in the installed version of SpamAssassin, but it is not easy to configure by the user. If you suspect an IP, you have two options to include the IP in the RBL list. The first one is updating the SpamAssassin to the latest version, since in every release of SpamAssassin includes the latest RBL. The second one is adding a firewall rule in IPTables. The second method is simple and easy too.

- ▶ `use_razor2 ( 0 | 1 )` (default: 1)

Whether to use Razor version 2, if it is available.

- ▶ `use_dcc ( 0 | 1 )` (default: 1)

Whether to use DCC, if it is available.

- ▶ `use_pyzor ( 0 | 1 )` (default: 1)

Whether to use Pyzor, if it is available.

## Language options

The Language options deal with the language codes and country codes for incoming mails.

- ▶ `ok_languages xx [ yy zz ... ]` (default: all)

This option is used to specify which languages are considered OK for incoming mail. SpamAssassin will try to detect the language used in the message text. Note that the language cannot always be recognized with sufficient confidence. In that case, no points will be assigned.

In your configuration, you must use the two or three letter language specifier in lowercase, not the English name for the language. You may also specify all if a desired language is not listed, or if you want to allow any language. The default setting is all.

Examples:

**ok\_languages all** Allow all languages

**ok\_languages en** Only allow English

**ok\_languages en ja ko** Allow English, Japanese, and Korean

**Attention:** If there are multiple `ok_languages` lines, only the last one is used.

- ▶ `ok_locales xx [ yy zz ... ]` (default: all)

This option is used to specify which locales (country codes) are considered OK for incoming mail. Mail using character sets used by languages in these countries will not be marked as possibly being spam in a foreign language.

If you receive lots of spam in foreign languages, and never get any non-spam in these languages, this may help. Note that all ISO-8859-\* character sets, and Windows code page character sets, are always permitted by default.

Set this to all to allow all character sets. This is the default.

Examples:

**ok\_locales all** Allow all locales  
**ok\_locales en** Only allow English  
**ok\_locales en ja ko** Allow English, Japanese and Korean

**Attention:** If there are multiple `ok_locales` lines, only the last one is used.

### 4.4.3 Configuring SpamAssassin for Postfix

SpamAssassin requires the Postfix `master.cf` configuration file to point to it, in addition to requiring its own configuration file, `/etc/mail/spamassassin/local.cf`, to be edited for proper functionality.

1. Edit the Postfix config file.

a. With the virtual console still open, at the command line type:

```
vi /etc/postfix/master.cf
```

b. Under `services`, locate the line that reads:

```
smtp inet n - n - - smtpd
```

and alter so that it reads:

```
smtp inet n - n - - smtpd -o content_filter=spamassassin
```

c. At the end of the file, add a line that reads:

```
spamassassin unix - n n - - pipe user=nobody argv=/usr/bin/spamc -f -e
/usr/sbin/sendmail -oi -f ${sender}${recipient}
```

making sure that it is all contained on one line.

d. Save the file and quit the editor.

2. Edit the SpamAssassin config file.

a. The configuration file is located at `/etc/mail/spamassassin/local.cf`. At the command line type:

```
vi /etc/mail/spamassassin/local.cf
```

b. Remove all existing lines in the file.

c. Add the following lines to the file:

```
required_hits 5.0
rewrite_subject 1
subject_tag [SPAM Detected]
report_safe 0
use_bayes 1
bayes_auto_learn 1
skip_rbl_checks 0
use_razor2 1
use_dcc 1
use_pyzor 1
ok_languages en
ok_locales en
```

- d. If you are experiencing bandwidth issues, or you receive a high volume of messages and mail delivery appears slow, try changing the following lines:

```
skip_rbl_checks 1
use_razor2 0
use_dcc 0
use_pyzor 0
```

This way, remote blacklisting will not be used, and verifying mail checksums remotely will not happen.

- e. Save the file and quit the editor.
3. Add the SpamAssassin process to system startup.
    - a. If you are configuring SpamAssassin on SLES9, at the command line type:

```
vi /etc/rc.d/rc
```

add a line at the end of file right before the line exit 0, which reads:

```
./usr/bin/spamd -d
```

Save the file and quit the editor.

- b. If you are configuring SpamAssassin on RHEL4, at the command line type:

```
vi /etc/rc.d/rc.local
```

add a line in the beginning of the file, which reads:

```
./etc/rc.d/init.d/spamassassin
```

Save the file and quit the editor.

4. Restart the Postfix process by executing the commands:

```
postfix stop
postfix start
```

SpamAssassin is now setup to process all incoming e-mail messages. Messages are given a score based on the analysis that they are given by SpamAssassin. If a message obtains a score above 5.0, that message has the subject line altered to contain the words “[SPAM Detected]”. End users may setup a filter in their mail client to send these messages directly to their trash.

#### 4.4.4 Configuring SpamAssassin for qmail

Since we are going to use qmail-scanner to integrate SpamAssassin with qmail, doing minimum configuration on SpamAssassin is enough now. This section describes the required configuration on SpamAssassin for qmail.

1. Edit the SpamAssassin config file.

The configuration file is located at /etc/mail/spamassassin/local.cf. At the command line type:

```
vi /etc/mail/spamassassin/local.cf
```

2. Remove all existing lines in the file.
3. Add the following lines to the file:

```
required_hits 5.0
rewrite_subject 1
subject_tag [SPAM Detected]
report_safe 0
use_bayes 1
bayes_auto_learn 1
```

```

skip_rbl_checks 0
use_razor2 1
use_dcc 1
use_pyzor 1
ok_languages en
ok_locales en

```

4. If you are experiencing bandwidth issues, or you receive a high volume of messages and mail delivery appears slow, try changing the following lines:

```

skip_rbl_checks 1
use_razor2 0
use_dcc 0
use_pyzor 0

```

This way, remote blacklisting will not be used, and verifying mail checksums remotely will not happen.

5. Save the file and quit the editor.

#### 4.4.5 Adding SpamAssassin process to system startup

SpamAssassin should be started when system gets boot up. This is done by including the spamd daemon in the system startup scripts.

1. If you are configuring SpamAssassin on SLES9, at the command line type:

```
vi /etc/rc.d/rc
```

add a line at the end of file right before the line exit 0, which reads:

```
./usr/bin/spamd -d
```

Save the file and quit the editor.

2. If you are configuring SpamAssassin on RHEL4, at the command line type:

```
vi /etc/rc.d/rc.local
```

add a line in the beginning of the file, which reads:

```
./etc/rc.d/init.d/spamassassin
```

Save the file and quit the editor.

#### 4.4.6 Installation and configuration of qmail-scanner

The qmail-scanner is an add-on for qmail which enables qmail to scan all e-mails when it is used in conjunction with external virus scanners and spam filters. In our setup we integrate qmail-scanner with Clam Antivirus and SpamAssassin.

##### Installation of qmail-scanner

**Important:** Ensure Clam Antivirus and SpamAssassin are installed and operational before attempting to install qmail-scanner.

1. The qmail-scanner should be run by the user and group of qscand. Create the user and group by typing:

```

groupadd qscand
useradd -g qscand -s /bin/false -c "Qmail-Scanner" qscand

```

2. Change into the directory /usr/local/src by typing:

```
cd /usr/local/src
```

3. Download the source tarball from sourceforge.net by executing the command:

```
curl -O
http://easynews.dl.sourceforge.net/sourceforge/qmail-scanner/qmail-scanner-1.25.tgz
```

**Note:** As per <http://qmail-scanner.sourceforge.net>, qmail-scanner requires a qmail queue patch to be included in qmail. Since we have used netqmail for our qmail installation, the patch is already included.

4. Unpack qmail-scanner by extracting the tarball. At the command prompt type:

```
tar xzf qmail-scanner-1.25.tgz
```

5. Change into the qmail-scanner source directory by typing:

```
cd qmail-scanner-1.25
```

6. Run ./configure with required options.

```
./configure --admin postmaster --domain secureyouremails.com --local-domains
secureyouremails.com --sa-quarantine 0.1 --sa-delete 5.0 --sa-subject "[SPAM Detected]"
--skip-setuid-test
```

**Note:** You have to replace `secureyouremail.com` by your own domain name.

While configuring it would prompt for Continue? ([Y]/N) few times. Just press Enter as the default option is to continue.

This will now auto-detect the filter software installed on your system, and will generate a script specific to your system. If you don't see any errors reported, then the build is, probably, successful.

7. Run ./configure again with --install along with the options you chose, this will do the same as the previous line, but will also create the directory structure required, and install qmail-scanner-queue.pl.

```
./configure --admin postmaster --domain secureyouremails.com --local-domains
secureyouremails.com --sa-quarantine 0.1 --sa-delete 5.0 --sa-subject "[SPAM Detected]"
--skip-setuid-test --install
```

**Tip:** Use the up arrow key to get the previous executed command at the command line and append --install.

8. qmail-scanner is installed now. Exit the directory by typing:

```
cd ..
```

9. The Figure 4-24 on page 226 shows the steps involved in installing qmail-scanner.

```

LINUXST:~ # groupadd qscand
LINUXST:~ # useradd -g qscand -s /bin/false -c "Qmail-Scanner" qscand
LINUXST:~ # cd /usr/local/src
LINUXST:/usr/local/src # curl -O
http://easynews.dl.sourceforge.net/sourceforge/qmail-scanner/qmail-scanner-1.25.tgz
 % Total % Received % Xferd Average Speed Time Curr.
 Dload Upload Total Current Left Speed
100 0k 100 0k 0 0 0k 0 0:00:01 0:00:01 0:00:00 0k
LINUXST:/usr/local/src # tar zxf qmail-scanner-1.25.tgz
LINUXST:/usr/local/src # cd qmail-scanner-1.25/
LINUXST:/usr/local/src/qmail-scanner-1.25 # ./configure --admin postmaster --domain
secureyouremails.com --local-domains secureyouremails.com --sa-quarantine 0.1
--sa-delete 5.0 --sa-subject "[SPAM Detected]" --skip-setuid-test
Building Qmail-Scanner 1.25...

This script will search your system for the virus scanners it knows
about, and will ensure that all external programs
qmail-scanner-queue.pl uses are explicitly pathed for performance
reasons.

It will then generate qmail-scanner-queue.pl - it is up to you to install it
correctly.

Continue? ([Y]/N) (Just press Enter here)

The following binaries and scanners were found on your system:

mimeunpacker=/usr/local/bin/reformime

Content/Virus Scanners installed on your System

clamdscan=/usr/local/bin/clamdscan

Qmail-Scanner details.

log-details=syslog
log-crypto=0

<some output is skipped here>

If that looks correct, I will now generate qmail-scanner-queue.pl
for your system...
Continue? ([Y]/N) (Just press Enter here)

Finished. Please read README(.html) and then go over the script to
check paths/etc, and then install as you see fit.

<remaining output is skipped here>

LINUXST:/usr/local/src/qmail-scanner-1.25 # ./configure --admin postmaster --domain
secureyouremails.com --local-domains secureyouremails.com --sa-quarantine 0.1
--sa-delete 5.0 --sa-subject "[SPAM Detected]" --skip-setuid-test --install

<output is skipped here>

LINUXST:/usr/local/src/qmail-scanner-1.25 # cd ..
LINUXST:/usr/local/src #

```

Figure 4-24 Installing qmail-scanner

## Configuration of qmail-scanner

To configure the qmail-scanner:

### 1. Making qmail-scanner-queue.pl work

Due to setuid issues with Perl, you need to use a C program to wrap qmail-scanner-queue.pl. This program is included with qmail-scanner in the contrib directory of the unpacked source.

- a. Change the directory to the unpacked source by typing:

```
cd /usr/local/src/qmail-scanner-1.25
```

- b. Now execute the following commands to wrap qmail-scanner-queue:

```
cd contrib
make
```

- c. Copy the wrapped qmail-scanner-queue to qmail's binary path.

```
cp qmail-scanner-queue /var/qmail/bin/
```

- d. Now you have two new executables in /var/qmail/bin, qmail-scanner-queue.pl and qmail-scanner-queue. Change these file permission to make them executable. Type the following commands:

```
cd /var/qmail/bin
chmod 6755 qmail-scanner-queue
chmod 0755 qmail-scanner-queue.pl
chown qscand:qscand qmail-scanner-queue*
```

### 2. Enabling qmail-scanner-queue in qmail-smtpd

- a. Edit the qmail-smtpd's run file by typing:

```
vi /var/qmail/supervise/qmail-smtpd/run
```

- b. Uncomment the line which reads:

```
QMAILQUEUE="/var/qmail/bin/qmail-scanner-queue" export QMAILQUEUE
```

- c. Save the file and quit the editor.

### 3. Testing the qmail-scanner-queue.pl

Open qmail-scanner-queue.pl and remove the -T option from the shebang line (the very first line that has the perl binary path), then save the file and exit the editor.

Test it by executing the command

```
./qmail-scanner-queue -z
./qmail-scanner-queue -v
```

If the output shows version information for qmail-scanner, the wrapper is working. Also, it confirms qmail-scanner picked up all of the right scanners.

### 4. Changing ownership of /var/spool/qmailscan

Since qmail-scanner uses the path /var/spool/qmailscan to handle the filtered mails, the directory qmailscan should be owned by qscand. Change the ownership by executing the command:

```
chown -R qscand:qscand /var/spool/qmailscan/
```

qmail-scanner will not work if any of the files it accesses are not owned by qscand.

5. Now you need to tell qmail to use the scanner to process incoming mail. This is done by modifying the smtp tcp rule we created earlier. Edit the file /etc/tcp.smtp and change the entries so that it reads as shown below.

```
127.:allow,RELAYCLIENT="",QS_SPAMASSASSIN="yes",QMAILQUEUE="/var/qmail/bin/qmail-scanner-queue"
```

```
:allow,QS_SPAMASSASSIN="yes",QMAILQUEUE="/var/qmail/bin/qmail-scanner-queue"
```

This will scan all mail, both internal and external. If you feel this is not necessary to scan the internal mails you can leave the lines which has the local IP address as it is.

6. Reload the qmail rules file by typing:

```
qmailctl cdb
```

7. Restart the qmail to include the changes into effect:

```
qmailctl restart
```

Now you have Clam AntiVirus, SpamAssassin and qmail-scanner all working together.

When a messages come into the server, qmail-scanner takes the message and pipes it out to both Clam Anti-Virus and SpamAssassin. If the message contains a virus, Clam AV quarantines it a `/var/spool/qmailscan/quarantine` and then send a notification e-mail to whoever you specified in the qmail-scanner installation.

If the message does not contain a virus, it is then scanned by SpamAssassin. Depending on the score that SpamAssassin assigns to the message, SpamAssassin will either let the message go unaltered to its destination or it will tag the message as `[SPAM Detected]`.

If the message is tagged as `[SPAM Detected]`, it will still arrive at its destination, but with an altered subject that will signal to the recipient that this was tagged as SPAM.

Once the message is tagged, the recipient can then configure his/her mail client to deal with those tagged message in whatever manner she sees fit.



## Part 2

# ISV solutions

This part discusses ISV-based network security solutions.

Archived



## StoneGate Firewall and VPN for iSeries

This chapter reviews the advantages and disadvantages of using a commercial firewall over an open source solution for Linux.

This chapter then provides you with an overview of the necessary planning and preparation tasks needed to implement the StoneGate Firewall.

This chapter does not take you through the installation of the firewall engine or of the management center, as this is covered in the StoneGate Firewall documentation supplied with the software.

The final section of this chapter quickly reviews a couple of other scenarios where the StoneGate Firewall can improve security and accountability for an eServer i5 infrastructure.

## 5.1 Why commercial firewalls

First let us review the benefits and disadvantages of open source and commercial firewall solutions and then between integrated and external firewalls.

### 5.1.1 Open source verses commercial firewalls

Why would you want to invest in a commercial firewall when both the eServer i5 supported distributions SUSE Linux Enterprise Server 9 (SLES 9) and Red Hat Enterprise Linux (RELH 4) have this functionality built in?

Using one of these distributions you would need to make use of iptables and netfilter, both of which are command line based tools, and in larger firewall configurations can be very large and confusing to follow. ISV solutions usually provide more advanced user interfaces, and this is especially true with the GUI from the StoneGate firewall solution.

Hardening of the kernel and installing and keeping up to date all the security fixes is also your responsibility. This is a very important process to insure the security of your firewall environment and, if not done correctly, it will lead to you locking yourself out of the partition or, in the worse case scenario, leaving backdoors open for hackers.

You will also need to remove unnecessary servers included and installed by default by the distribution software such as DNS, DHCP, and HTTP.

All of these tasks take time and experience to perform and keep up to date, all of which add to the total cost of ownership of your firewall. An ISV solution will perform most, if not all, of these tasks for you, so reducing the administrative tasks needed to install and maintain your firewall.

### 5.1.2 Internal verses external firewalls

Why would you want to invest in a firewall installation on an eServer i5 system when you could just as easily make do with an external off-the-shelf appliance? Or even your existing external firewalls?

The answer to these questions very quickly becomes apparent when you start to look at the complexity, manageability, and performance issues you could have with an external firewall. If you have multiple partitions on an eServer i5 system, for example, and you have external firewalls, each partition would have to have dedicated physical network adapters so you could control access between partitions and users.

On top of this, if you have a multi-layered application infrastructure, this could lead to a situation where you have traffic traveling outside of the partition via the external firewall to the backend partition and then back again just to complete a single transaction.

To further complicate this solution, depending on your business requirements, you may need to consider a high availability firewall, either a cluster or hot standby.

This can be seen better in Figure 5-1 on page 233 just how complex this can get, and how much hardware could be required. This figure does not attempt to show any of the firewall security policies that would be required to secure this environment.




Figure 5-1 External firewall solution

## 5.2 The StoneGate Firewall

This section provides an overview of Stonesoft's StoneGate High Availability Firewall and VPN solution. We review the important features of the product and their benefits in improving security and availability for the eServer i5 infrastructure.

The StoneGate High Availability Firewall and VPN solution includes its own integrated and hardened operating system. It is built on the principle that everything is denied unless it is expressly permitted. The StoneGate Firewall not only provides a true stateful inspection firewall, but also multi-layer inspection. This means that the firewall can function as a packet filter, a stateful inspection firewall, or where required as an application-level firewall, all available on a rule-by-rule basis.

Communications between all components are encrypted by default; communications are also authenticated using advanced PKI digital certificates. SSH-based interaction with the firewalls is possible, but this is not enabled by default, and the security policy must permit this type of communications.

Stonesoft has built their firewall system around a distributed architecture allowing you to deploy the system components effectively to different network environments. This total system can be managed via a single user interface. The StoneGate system is built up of three components: The firewall engine, the StoneGate Management Center (SMC), and the GUI Client (Management Client).






Figure 5-3 Full Mesh topology

Within a Full Mesh topology all gateways connect to each other. This provides VPN connectivity directly between each site.




Figure 5-4 Star topology

A Star topology has direct connections to the selected central *non-spoke* gateway only.

It is possible to have a combination of both of these types bringing branch offices to a central hub and then connecting to a Full Mesh VPN network.

VPN configuration and management is done from within the VPN Manager (part of the SMC; see 5.2.5, “Centralized management center: SMC” on page 237). With the *build once use anywhere* components it is easy for administrators to implement highly secure, encrypted, and authenticated VPN connections.

The free VPN Client for Windows is provided by Stonesoft to facilitate VPN communications between remote clients and central systems. This client can be installed either with or without application security.

Without application security the security level is set automatically or user chosen. These levels allow all traffic, allow all outbound traffic, allow only VPN traffic, or allow no traffic at all. Additionally, the boot security feature prevents all inbound connections, if the VPN Client is exited or terminated. These settings will affect all traffic to and from the computer.

The application security adds a *simplified personal firewall*. This gives the benefit of allowing the user to control access by application.

VPN is an integral part of the product and is Internet Protocol Security (IPsec) compliant. IPsec is a set of protocols developed by the Internet Engineering Task Force (IETF), the main standards organization for the Internet to support secure exchange of packets at the IP layer.

### 5.2.3 Firewall configuration and maintenance

There is no need to physically access the firewalls for upgrades, as this can be done remotely through the GUI Client.

The firewall includes failsafe policy upload, so if the policy fails to install correctly for any reason the firewall will reinstate the old security policy, and failsafe remote upgrade, and if the upgrade fails the firewall can be reverted to the previous level.

The firewall sends performance statistics along with state information and log data to the SMC and receives configuration changes.

Configuration and policy updates are sent to the firewall from the Management Server. This configuration information is by default encrypted on the system.

### 5.2.4 StoneGate management client

All configuration and management tasks are performed via the Administration (GUI) Client:

- ▶ Firewall and VPN configuration
- ▶ Creating network elements (for example, servers, firewalls, VPNs, hosts, and services), the building blocks for describing networks and systems
- ▶ Designing security and alert policies
- ▶ Monitoring system operation and log data
- ▶ Remotely upgrading the firewall engines

This interface is the most visible part of the system to the administrator.

It is possible to deploy the GUI client anywhere within the network, and for multiple clients to run at the same time. This client manages both the StoneGate Firewall/VPN products and the StoneGate IPS, improving the manageability of the entire network security infrastructure.

**Note:** The GUI never communicates directly with the firewall engines. All firewall configuration changes are intermediated by the Management Center for security reasons.




Figure 5-5 Administration Client Control Panel

Figure 5-5 illustrates a sample screen of the management GUI client.

### 5.2.5 Centralized management center: SMC

The StoneGate Management Center (SMC) allows you, with the same graphical user interface, to manage both the StoneGate Firewalls/VPNs and the StoneGate IPS. Also, logs and alerts generated by both of these products can be easily handled efficiently from within the SMC.

The SMC comprises two components:

- ▶ The Management Server
  - One per firewall environment
- ▶ The Log Server
  - One or more per firewall environment

You can install these servers on the same system or you may distribute them on separate systems or locations.

## Management Server

This is the central point of administration for a StoneGate Firewall environment, storing all configuration information, and firewall and network configuration and security policy information.

## Log Server

As implied by the name, Log Servers are used to store and manage the logs and alerts produced by the firewall. They also provide the GUI client with log and alert information as required.

Not only do you have the option to run the log server on the same physical server or separate servers as the management server, it is also possible to have multiple firewalls logging to one Log Server. You could also use a Log Server for logs and forward all alerts to a separate Log Server. This can make alert escalation simpler to maintain.

If the log server is down for any reason then the logging information is held on the firewall until the log server comes back online.

**Note:** During installation, the Management Server and Log Server can be configured as services for starting and stopping at the same time as Windows.

## 5.2.6 Reporting tools

The SMC also provides extensive reporting abilities; with the reporting tools it is possible to generate statistical reports on the logged events. The SMC comes with a wide range of default reports that can be customized to your requirements. These can be displayed on screen or exported as plain text reports for further processing or in PDF format for printing or displaying on the Web or distribution via e-mail, as illustrated in Figure 5-6 on page 239.




Figure 5-6 SMC Reporting Manager

## 5.2.7 Clustering and VPN high availability

Two more important features of the StoneGate VPN firewall solution are Multi-Link VPN and Clustering.

### Multi-Link VPN

With the StoneGate Multi-Link technology it is possible to load balance VPN traffic between multiple network links (leased lines, dialup modems, or ISPs). By using multiple links you also increase the availability of communications and minimize the effects of dropped connections, as illustrated in Figure 5-7 on page 240.




Figure 5-7 Multi-Link VPN with single security gateways

### Clustering and high availability

Although clustering is not available for an eServer i5 firewall solution it is important to mention this here as a major advantage of the product.

Multiple firewalls can be installed to form a cluster that functions as a single entity. A StoneGate cluster can accommodate up to 16 firewall engines. This clustering technology, as illustrated in Figure 5-8 on page 241, provides load balancing and high availability to the StoneGate environment, which in turn provides high performance and reliability.




Figure 5-8 Example of StoneGate cluster architecture

For a further understanding of how this can be accomplished, see 5.4.2, “External Firewall Cluster and integration environment” on page 244.

### 5.3 Planning and preparation for the StoneGate Firewall

This section helps you prepare for a StoneGate Firewall installation in an eServer i5 environment. We review with you the hardware and software requirements and then look at sizing and performance tips.

**Note:** The eServer i5 Service Processor firmware needs to be at GA2 level or later, and the HMC needs to be at Version 4 Release 2.0 level 20040816 or later.

### 5.3.1 Hardware and software requirements for eServer i5 and SMC

In this section we provide you with the minimum requirements for the firewall LPAR on the eServer i5 and the SMC server.

#### StoneGate Firewall Version 2.2.9

This is the first General Release (GA) version to run on the IBM eServer i5 system.

This version is compatible with:

- ▶ IBM POWER5 processors
- ▶ i5/OS Version 5 Release 3

We recommend using the following StoneGate component versions:

- ▶ StoneGate Management System v3.1.2
- ▶ StoneGate VPN Client v2.5.2
- ▶ StoneGate Server Pool Monitoring Agent v2.2.0

**Note:** The IBM DynamicRM and ServiceRM are not implemented in this version of StoneGate. Because of this, shutting down the StoneGate LPAR via the HMC is done forcefully. The `halt` command should be used from within the StoneGate console, before shutting down the LPAR via the HMC.

#### Firewall partition requirements

The minimum requirements for each StoneGate Firewall instance on an eServer i5 are listed below.

- ▶ Dedicated StoneGate logical partition on the eServer i5
- ▶ Minimum for 0.1 CPU allocation
- ▶ Minimum of 256 MB main storage
- ▶ Minimum of 2 GB storage space

#### StoneGate Management Center requirements

The basic hardware requirements for the SMC are:

- ▶ Pentium® III processor or later recommended (suggested minimum processor speed 1 GHz) or equivalent on a non-Intel platform
- ▶ Mouse or pointing device required for GUI installations
- ▶ SVGA (1024x768) display or higher required for GUI installations
- ▶ 512 MB RAM
- ▶ Disk space for Management Server: 2 GB
- ▶ Disk space for Log Server: 20 GB–80 GB

Operating systems and versions that the SMC is supported on:

- ▶ Microsoft Windows 2003
- ▶ Microsoft Windows XP SP1 and SP2
- ▶ Microsoft Windows 2000 SP3 and SP4
- ▶ Red Hat Linux 7.3 and 9
- ▶ Red Hat Enterprise Linux 2.1 and 3.0
- ▶ Sun™ Solaris™ 8 and 9

**Note:** At the time of writing Stonesoft had only tested the U.S. English versions of the Windows operating systems, but it is believed that other locales will work as well.

## 5.4 Other scenarios

In this final section we provide you with information about how the StoneGate Firewall could be extended.

You will see that it is not only possible to protect LPARs, eServer i5/OS, Linux, or AIX, but also integrated Windows servers. This could even be extended to include servers that have not yet been integrated or cannot be integrated into an eServer i5 solution because of technical or business reasons.

You will see that it also possible with a StoneGate Firewall solution to extend your manageability to cover your total enterprise infrastructure, with a combination of integrated and appliance firewalls, benefiting from the unique StoneGate Management Center.

### 5.4.1 Integration environment

Let us say that we have a firewall to protect our mail server from both the Internet and from our internal end users.

First we quickly look at how it is possible to extend this to include integrated Window servers and standalone servers.

Normally when you integrate Windows servers into an eServer i5 environment you use the built-in Ethernet adapters to provide the end users with connection to these servers. In this scenario we do not use these adapters, as with our mail environment all communications go through the firewall partition.

For external servers this would only require one or more physical Ethernet adapters being added to the eServer i5 and then allocating them to the firewall partition.

There is only one major issue that would need to be taken into account for this extension to our environment and that is the size and amount of data that you could or would want to practically send over this connection. As long as the end users are not pulling gigabytes of data from standalone servers this solution should have no performance problems. Figure 5-9 on page 244 illustrates this scenario.




Figure 5-9 Integration environment

#### 5.4.2 External Firewall Cluster and integration environment

Because of the way that Stonesoft has developed their firewall it is possible to manage all your StoneGate Firewalls from one Management Center, which brings both improvements in manageability and savings in cost.

For this scenario what you would need to do is create an appliance firewall or cluster, depending on the requirements, with an integrated firewall on the eServer i5 system or systems.

With this solution you would be able to have multiple ISPs to provide high availability for your Internet access, and if you had more than one IBM eServer i5 then you could provide load balancing for your Web applications. Figure 5-10 on page 245 illustrates this scenario.




Figure 5-10 Integration environment with cluster

Archived

## MPP Service Provider Edition: Foundation of E-mail Filtering Services

The Message Processing Platform (MPP) is an e-mail security middleware application that is used to build scalable and flexible e-mail filtering, archival, and compliance services. MPP replaces complex, internally developed filtering applications; inflexible antivirus/antispam appliances; point antivirus and antispam applications; and open source filtering tools with a scalable policy-based application framework that allows for easy creation of a completely custom messaging architecture. MPP unifies the distinct elements of a filtering service including iSeries e-mail servers, commercial and open-source scanning engines, provisioning systems, management tools, security controls, quarantine review services, and more into a cohesive, easy-to-manage application. MPP provides the integration point and a single point of management, monitoring, and configuration for the key service elements of filtering services.

There are many benefits when using MPP as the basis for e-mail security applications:

- ▶ Fast response to new business requirements including implementing new spam and virus technologies, e-mail archival and search, regulatory compliance, and content filtering
- ▶ Simple integration of open source and commercial scan engines
- ▶ Reduced time-to-market for new services
- ▶ Lower operational costs of filtering services through unified management, configuration, and control of diverse service elements
- ▶ A completely customizable application framework that can be tailored to reflect the service provider's brand and business requirements
- ▶ A single Web-based quarantine for all end users irrespective of the spam and virus engines chosen and that does not change when engines are swapped

## 6.1 The core of your e-mail service: Policy-based service offerings

MPP sits in the core of your message filtering service analyzing traffic, processing messages based on the service policy configuration for the e-mail, and keeping detailed statistics.

The policy engine is the central component of MPP and it is what makes differentiated service offerings possible with MPP. There are three key elements to the policy engine: The classifier, the service policy configurations, and the policy membership databases. The service policy configuration is predefined in the MPP XML configuration file and contains the complete service configuration for all elements of the policy including options such as which scanners to use, which security checks to invoke, whether to archive, whether to accept or reject the e-mail based on white or black list information, whether to filter the e-mail for content, what type of content to filter, and so on.

The policy membership database contains a list of e-mail addresses and a policy name associated with the address. The policy membership database may be integrated with an existing LDAP directory, stored in text files, or integrated directly into the configuration file. The MPP classifier decodes SMTP envelope information for each e-mail and finds policy matches for the e-mail senders and receivers by performing lookups in the policy membership database. After a match is found the classifier then processes the e-mail based on the predefined service policy configuration, or the default policy if no match is found.

## 6.2 Selected MPP-based service offerings

MPP-based service offerings include:

- ▶ Tiered, opt-in e-mail antivirus and antispam offerings. A low-level service could check for viruses only, while the high-level service will apply multiple scanners, discard all viruses, and quarantine spam for end-user review.
- ▶ E-mail archival and cataloging. Offer an e-mail archival service based on standard databases or text files. Utilize standard search tools for fast search and retrieval.
- ▶ Restricted communication services. Allow your clients to restrict which senders can send to the Internet or explicitly define the allowed senders or receivers on a per-domain or per-address basis.
- ▶ Provide a branded interface for end users to review quarantined spam and archived mail. Change the interface on a per-client basis.
- ▶ Offer your client detailed graphical statistics on their e-mail usage using our open SNMP engine.
- ▶ Provide attachment filtering and archival services.
- ▶ Offer any or all of the above application ideas simultaneously to any arbitrary set of clients.

## 6.3 MPP engine and e-mail server support

The e-mail server support includes:

- ▶ Sendmail
- ▶ Qmail
- ▶ Postfix
- ▶ CommuniGate Pro
- ▶ Open-Xchange

The iSeries engine support includes:

- ▶ Sophos
- ▶ Cloudmark
- ▶ PatternAuthority
- ▶ SpamAssassin
- ▶ ClamAV

## 6.4 Managing MPP

MPP integrates with Webmin, an open source tool to manage all aspects of a Linux server. The MPP Webmin module allows for configuration of all scanning engines, quarantine control, log monitoring, virus and spam updates, and more.

## 6.5 Conclusion

MPP is an e-mail security application that is designed to be the core of your e-mail filtering application. The MPP has many advantages over the common approaches to building a filtering service such as re-inventing the wheel with home-grown software; using open source integration tools; using closed, proprietary e-mail filtering appliances; or using point software solutions. MPP provides the best of all solutions by enabling rapid deployment of a wide range of e-mail applications that can be completely customized to suit your business needs. MPP integrates a wide range of open source, embedded, and commercial application elements into a cohesive e-mail security application.

If you need to contact Message Partners Inc., the provider of MPP, please use the following contact information:

Message Partners Inc.  
271 North Ave Suite 1210  
New Rochelle, NY 10801  
(877) 302-2027  
+1 (914) 712-9050  
(914) 206-9609 – Fax

<mailto:info@messagepartners.com>  
<http://www.messagepartners.com>

Archived



## Bytware's StandGuard Anti-Virus for Linux

Systems running Linux, including the iSeries running Linux in a partition, can store viruses and spread them to other systems and platforms through file access.

With more than 100,000 known viruses floating around today, new virus threats occurring weekly, and evidence that new viruses are being designed to target corporate enterprises, virus detection for any operating system is more important now than ever before. Today it is no longer simply a virus threat; Internet worms, Trojan horses, and backdoors are now a significant and growing threat, alongside .EXE infectors and macro viruses. These blended threats, also known as *malware* for malicious software, are the viruses of the new millennium.

Threats can come from anywhere, and once on your server, your network is in danger. You should scan all servers to ensure that an embedded virus, Trojan horse, or worm is not going to wreak havoc.

While many companies rely on a freeware product, this may not be sufficient for your company's needs. Freeware offers a cost advantage, but lacks the testing and support that many companies desire—especially when it comes to system security and legal requirements like those mandated by legislation such as the Sarbanes-Oxley Act.

StandGuard Anti-Virus for Linux by Bytware is a variation of the award-winning StandGuard Anti-Virus for iSeries, introduced as the only native AV solution for iSeries and eServer i5 systems in June of 2003.

StandGuard Anti-Virus for Linux is the result of collaboration between IBM, McAfee, and Bytware, with the resulting solution being based on the McAfee scanning engine. StandGuard Anti-Virus for Linux is a virus detection solution for systems running Linux, including the iSeries running Linux in a partition. StandGuard Anti-Virus for Linux is designed to scan files for viruses and perform advanced cleaning and notification procedures. StandGuard Anti-Virus for Linux incorporates the latest generation of McAfee's scanning engine, in turn making StandGuard Anti-Virus for Linux a mature product backed by battle-tested technology, advanced heuristic analysis, and generic detection and cleaning.

## 7.1 The McAfee virus scanning engine

The McAfee virus scanning engine:

- ▶ Scans within compressed files
- ▶ Decompresses and scans files compressed in packages (for example, PKZip, .LHA, .ARJ)
- ▶ Detects and cleans macro and script viruses
- ▶ Detects and cleans encrypted and polymorphic viruses
- ▶ Detects and cleans new viruses in executable files and OLE compound documents
- ▶ Detects and removes Trojan horses, worms, and many other types of malicious software (malware)
- ▶ Upgrades easily to new anti-virus technology

## 7.2 Additional StandGuard Anti-Virus features

This section covers the additional features of StandGuard Anti-Virus.

### 7.2.1 Command-line interface and graphical user interface

Whether you use the command line interface or the graphical interface, you will find StandGuard Anti-Virus for Linux simple and flexible to use. StandGuard Anti-Virus for Linux provides low-level commands that you can embed into your applications or nightly procedures, as illustrated in Figure 7-1.



```
standguard@tp600x:/home/standguard
File Edit View Terminal Go Help
[standguard@tp600x standguard]$ su root
Password:
[root@tp600x standguard]# avcfg -pmg -j "/usr/bin" -s now "All user bin"
avcfg:when task is completed output will be in:/home/standguard/log/All user bin
.log
[root@tp600x standguard]#
```

Figure 7-1 Command-line user interface

### 7.2.2 Automatic download of virus definitions (.DAT files)

StandGuard Anti-Virus for Linux ensures that you always have the latest protection against current virus threats by automatically downloading virus definition files from McAfee. By keeping the virus definition files up-to-date automatically, as illustrated in Figure 7-2 on page 253, StandGuard Anti-Virus for Linux protects you from the numerous new virus threats that occur.




Figure 7-2 Updating virus definitions

### 7.2.3 Automatic download of software updates and fixes

StandGuard Anti-Virus for Linux keeps itself up-to-date by downloading new features, fixes, and enhancements from Bytware.

### 7.2.4 Built-in scheduling features for scanning and updating

Built on Bytware's proven experience with administration, security, and automation, StandGuard Anti-Virus for Linux was designed from the ground up as a secure, automated Anti-Virus solution that prevents headaches, and does not create new ones. StandGuard Anti-Virus for Linux integrates with CRON to schedule automatic updating of virus definitions, product enhancements, and scanning tasks that you create. By automating these tasks using the iSeries tools that you know and trust, you can rest assured that StandGuard Anti-Virus for Linux is providing reliable, around-the-clock protection.




Figure 7-3 Built-in scheduling features

## 7.2.5 Network enabled

StandGuard Anti-Virus for Linux can retrieve virus definitions and program updates from either an FTP server or a shared local network path. The path can be located on any network path of your choice. This allows you to use one server or servers to download the virus definitions (from McAfee's FTP server), and the remaining servers or partitions can retrieve their virus definition files from the network, as illustrated in Figure 7-4 on page 255.




Figure 7-4 Network-enabled program updates

## 7.2.6 Extensive logging capabilities

As important as it is to install anti-virus protection on your server, it is equally important to know when problems occur. Two important events that you need to monitor are:

- ▶ StandGuard Anti-Virus for Linux detected and removed virus.
- ▶ If virus definition files could not be retrieved. In addition, you could monitor other events, such as if a scan ended abnormally or did not run at all.

StandGuard Anti-Virus for Linux will log all scan events to a file for easy reporting, as illustrated in Figure 7-5 on page 256.




Figure 7-5 Viewing the reports

For additional information about StandGuard Anti-Virus for Linux visit:

<http://www.bytware.com/products/sgav.html>



**A**

## **Tips and techniques**

This section contains the planning worksheets used throughout this book. It also includes other useful information about and references to e-mail security tools we used. We suggest that these be printed off and completed for future use.

# SpamAssassin: Language options

SpamAssassin supports many languages and locales. The following tables list the available language and locale options.

## ok\_language

Usage: `ok_language <the_two_letter_language_identifier>`

Example: `ok_language hi ta`

Table A-1 List of supported languages in SpamAssassin

|                         |                   |                 |                         |                 |
|-------------------------|-------------------|-----------------|-------------------------|-----------------|
| af - Afrikaans          | en - English      | hu - Hungarian  | ne - Nepali             | sq - Albanian   |
| am - Amharic            | eo - Esperanto    | hy - Armenian   | nl - Dutch              | sr - Serbian    |
| ar - Arabic             | es - Spanish      | id - Indonesian | no - Norwegian          | sv - Swedish    |
| be - Byelorussian       | et - Estonian     | is - Icelandic  | pl - Polish             | sw - Swahili    |
| bg - Bulgarian          | eu - Basque       | it - Italian    | pt - Portuguese         | ta - Tamil      |
| bs - Bosnian            | fa - Persian      | ja - Japanese   | qu - Quechua            | th - Thai       |
| ca - Catalan            | fi - Finnish      | ka - Georgian   | rm - Rhaeto<br>-Romance | tl - Tagalog    |
| cs - Czech              | fr - French       | ko - Korean     | ro - Romanian           | tr - Turkish    |
| cy - Welsh              | fy - Frisian      | la - Latin      | ru - Russian            | uk - Ukrainian  |
| da - Danish             | ga - Irish Gaelic | lt - Lithuanian | sa - Sanskrit           | vi - Vietnamese |
| de - German             | he - Hebrew       | lv - Latvian    | sco - Scots             | yi - Yiddish    |
| el - Greek              | hi - Hindi        | mr - Marathi    | sk - Slovak             | zh - Chinese    |
| gd - Scottish<br>Gaelic | hr - Croatian     | ms - Malay      | sl - Slovenian          |                 |

## ok\_locales

Usage: `ok_locales <the_two_letter_locale_identifier>`

Example: `ok_locales ko`

Table A-2 List of available locale options in SpamAssassin

|                                                               |
|---------------------------------------------------------------|
| en - Western character sets in general                        |
| ja - Japanese character sets                                  |
| ko - Korean character sets                                    |
| ru - Cyrillic character sets                                  |
| th - Thai character sets                                      |
| zh - Chinese (both simplified and traditional) character sets |

## qmail control files

Unlike other packages, qmail uses one file each for every variable, with the variable name as the file name, and the content in the file as the value for the variable. These files are called control files since they are placed in the qmail configuration control folder located at `/var/qmail/control`. qmail has 30 such files. Not all files are required. They are listed in the following table.

Table A-3 qmail control files

| Control             | Default value  | Used by                  | Purpose                                     |
|---------------------|----------------|--------------------------|---------------------------------------------|
| badmailfrom         | none           | qmail-smtpd              | blacklisted From addresses                  |
| bouncefrom          | MAILER-DAEMON  | qmail-send               | username of bounce sender                   |
| bouncehost          | me             | qmail-send               | host name of bounce sender                  |
| concurrencyincoming | none           | /service/qmail-smtpd/run | max simultaneous incoming SMTP connections  |
| concurrencylocal    | 10             | qmail-send               | max simultaneous local deliveries           |
| concurrencyremote   | 20             | qmail-send               | max simultaneous remote deliveries          |
| defaultdelivery     | none           | /var/qmail/rc            | default .qmail file                         |
| defaultdomain       | me             | qmail-inject             | default domain name                         |
| defaulthost         | me             | qmail-inject             | default host name                           |
| databytes           | 0              | qmail-smtpd              | max number of bytes in message (0=no limit) |
| doublebouncehost    | me             | qmail-send               | host name of double bounce sender           |
| doublebounceto      | postmaster     | qmail-send               | user to receive double bounces              |
| envnoathost         | me             | qmail-send               | default domain for addresses without "@"    |
| helohost            | me             | qmail-remote             | host name used in SMTP HELO command         |
| idhost              | me             | qmail-inject             | host name for Message-ID's                  |
| localiphost         | me             | qmail-smtpd              | name substituted for local IP address       |
| locals              | me             | qmail-send               | domains that we deliver locally             |
| me                  | FQDN of system | various                  | default for many control files              |
| morercpthosts       | none           | qmail-smtpd              | secondary rcpthosts database                |
| percenthack         | none           | qmail-send               | domains that can use "%"-style relaying     |

| Control        | Default value | Used by      | Purpose                                           |
|----------------|---------------|--------------|---------------------------------------------------|
| plusdomain     | me            | qmail-inject | domain substituted for trailing "+"               |
| qmqpservers    | none          | qmail-qmqpc  | IP addresses of QMQP servers                      |
| queuelifetime  | 604800        | qmail-send   | seconds a message can remain in queue             |
| rcpthosts      | none          | qmail-smtpd  | domains that we accept mail for                   |
| smtpgreeting   | me            | qmail-smtpd  | SMTP greeting message                             |
| smtproutes     | none          | qmail-remote | artificial SMTP routes                            |
| timeoutconnect | 60            | qmail-remote | how long, in seconds, to wait for SMTP connection |
| timeoutremote  | 1200          | qmail-remote | how long, in seconds, to wait for remote server   |
| timeoutsmtpd   | 1200          | qmail-smtpd  | how long, in seconds, to wait for SMTP client     |
| virtualdomains | none          | qmail-send   | virtual domains and users                         |

## Packages: Links to download

In this redpaper we covered the installation of many open source packages, downloading them from their mirror sites. In our examples, we have given direct links for where to download the packages. If any of the links do not work, it means either the package has been replaced by a new version or the mirror site is not available from your location for the time being. If you have any problem downloading those packages using the given URLs, please refer the relevant Web site listed here to get the valid link for downloading. Then replace the URL with our example.

- ▶ Postfix:  
<http://www.postfix.org/download.html>
- ▶ netqmail:  
<http://www.qmail.org>
- ▶ uspci-tcp:  
<http://cr.yip.to/software.html>
- ▶ daemontools:  
<http://cr.yip.to/software.html>
- ▶ fastforward:  
<http://cr.yip.to/software.html>
- ▶ maildrop:  
<http://www.courier-mta.org/download.php#maildrop>
- ▶ zlib:  
<http://www.zlib.net>

- ▶ GMP:  
<http://www.swox.com/gmp/#DOWNLOAD>
- ▶ Clam AntiVirus:  
<http://www.clamav.net/stable.php#pagestart>
- ▶ Clam SMTP:  
<http://memberwebs.com/nielsen/software/clamsmtp>
- ▶ Razor:  
<http://razor.sourceforge.net>
- ▶ DCC:  
<http://www.dcc-servers.net/dcc/source>
- ▶ Pyzor:  
<http://pyzor.sourceforge.net>
- ▶ qmail-scanner:  
<http://qmail-scanner.sourceforge.net>

Archived

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

## IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 264. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *Domino 6 for iSeries Best Practices Guide*, SG24-6937
- ▶ *IBM PowerPC: A guide to working with LPAR on Power5 IBM eServer i5 servers*, SG24-8000

## Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ Bynari:  
<http://www.bynari.net/>
- ▶ Postfix:  
<http://www.postfix.org/download.html>
- ▶ netqmail:  
<http://www.qmail.org>
- ▶ usspi-tcp:  
<http://cr.yp.to/software.html>
- ▶ daemontools:  
<http://cr.yp.to/software.html>
- ▶ fastforward:  
<http://cr.yp.to/software.html>
- ▶ maildrop:  
<http://www.courier-mta.org/download.php#maildrop>
- ▶ zlib:  
<http://www.zlib.net>
- ▶ GMP:  
<http://www.swox.com/gmp/#DOWNLOAD>
- ▶ Clam AntiVirus:  
<http://www.clamav.net/stable.php#pagestart>
- ▶ Clam SMTP:  
<http://memberwebs.com/nielsen/software/clamsmtp>

- ▶ Razor:  
<http://razor.sourceforge.net>
- ▶ DCC:  
<http://www.dcc-servers.net/dcc/source>
- ▶ Pyzor:  
<http://pyzor.sourceforge.net>
- ▶ qmail-scanner:  
<http://qmail-scanner.sourceforge.net>
- ▶ For SLES9 users, refer to the following link for information about the support from Novell:  
[http://support.novell.com/linux/linux\\_server\\_support.html](http://support.novell.com/linux/linux_server_support.html)
- ▶ For RHEL4 users, refer to the following link for information about the support from Red Hat:  
<http://www.redhat.com/software/rhel/compare/server/>
- ▶ IBM Global Services (IGS) via SupportLine:  
<http://www.ibm.com/services/us/index.wss/offering/its/a1000030>
- ▶ PuTTY download:  
<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

## How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)





# Secure Your E-mail Server on IBM *e*server i5 with Linux



**Understanding security issues for network and e-mail server**

**Linux open source solutions to secure your e-mail server**

**Linux-based ISV solutions to secure your e-mail server**

This IBM Redpaper will help you design a solution to protect your e-mail servers on IBM eServer i5 using various security solutions of Linux. Like many of Linux solutions, security implementation can be done using either free software solutions (that is, Open Source Software based solutions) or commerce software solutions (that is, Independent Software Vendor based solutions).

This redpaper has two main parts: Part I covers OSS-based solutions and Part II covers ISV-based solutions. In either case, the goal protecting your e-mail servers from various attacks.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

### **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)