

Fedora 12

security-guide

A Guide to Securing Fedora Linux



Johnray Fuller

John Ha

David O'Brien

Scott Radvan

Eric Christensen

Adam Ligas

Fedora 12 security-guide

A Guide to Securing Fedora Linux

Edition 12.1

Author	Johnray Fuller	jrfuller@redhat.com
Author	John Ha	jha@redhat.com
Author	David O'Brien	daobrien@redhat.com
Author	Scott Radvan	sradvan@redhat.com
Author	Eric Christensen	sparks@fedoraproject.org
Author	Adam Ligas	adam@physco.com

Copyright © 2009 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. The original authors of this document, and Red Hat, designate the Fedora Project as the "Attribution Party" for purposes of CC-BY-SA. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

For guidelines on the permitted uses of the Fedora trademarks, refer to https://fedoraproject.org/wiki/Legal:Trademark_guidelines.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

All other trademarks are the property of their respective owners.

The Fedora Security Guide is designed to assist users of Fedora in learning the processes and practices of securing workstations and servers against local and remote intrusion, exploitation, and malicious activity.

Focused on Fedora Linux but detailing concepts and techniques valid for all Linux systems, The Fedora Security Guide details the planning and the tools involved in creating a secured computing environment for the data center, workplace, and home.

With proper administrative knowledge, vigilance, and tools, systems running Linux can be both fully functional and secured from most common intrusion and exploit methods.

Preface	vii
1. Document Conventions	vii
1.1. Typographic Conventions	vii
1.2. Pull-quote Conventions	viii
1.3. Notes and Warnings	ix
2. We Need Feedback!	x
1. Security Overview	1
1.1. Introduction to Security	1
1.1.1. What is Computer Security?	1
1.1.2. SELinux	3
1.1.3. Security Controls	3
1.1.4. Conclusion	4
1.2. Vulnerability Assessment	5
1.2.1. Thinking Like the Enemy	5
1.2.2. Defining Assessment and Testing	6
1.2.3. Evaluating the Tools	7
1.3. Attackers and Vulnerabilities	9
1.3.1. A Quick History of Hackers	9
1.3.2. Threats to Network Security	10
1.3.3. Threats to Server Security	11
1.3.4. Threats to Workstation and Home PC Security	13
1.4. Common Exploits and Attacks	13
1.5. Security Updates	16
1.5.1. Updating Packages	16
1.5.2. Verifying Signed Packages	16
1.5.3. Installing Signed Packages	17
1.5.4. Applying the Changes	18
2. Securing Your Network	21
2.1. Local users may install trusted packages	21
2.2. Workstation Security	21
2.2.1. Evaluating Workstation Security	21
2.2.2. BIOS and Boot Loader Security	22
2.2.3. Password Security	24
2.2.4. Administrative Controls	29
2.2.5. Available Network Services	36
2.2.6. Personal Firewalls	39
2.2.7. Security Enhanced Communication Tools	39
2.3. Server Security	40
2.3.1. Securing Services With TCP Wrappers and xinetd	40
2.3.2. Securing Portmap	43
2.3.3. Securing NIS	44
2.3.4. Securing NFS	46
2.3.5. Securing the Apache HTTP Server	47
2.3.6. Securing FTP	48
2.3.7. Securing Sendmail	50
2.3.8. Verifying Which Ports Are Listening	51
2.4. Single Sign-on (SSO)	53
2.4.1. Introduction	53
2.4.2. Getting Started with your new Smart Card	54
2.4.3. How Smart Card Enrollment Works	56

- 2.4.4. How Smart Card Login Works 56
- 2.4.5. Configuring Firefox to use Kerberos for SSO 57
- 2.5. Pluggable Authentication Modules (PAM) 59
 - 2.5.1. Advantages of PAM 60
 - 2.5.2. PAM Configuration Files 60
 - 2.5.3. PAM Configuration File Format 60
 - 2.5.4. Sample PAM Configuration Files 63
 - 2.5.5. Creating PAM Modules 64
 - 2.5.6. PAM and Administrative Credential Caching 64
 - 2.5.7. PAM and Device Ownership 66
 - 2.5.8. Additional Resources 67
- 2.6. TCP Wrappers and xinetd 68
 - 2.6.1. TCP Wrappers 69
 - 2.6.2. TCP Wrappers Configuration Files 70
 - 2.6.3. xinetd 77
 - 2.6.4. xinetd Configuration Files 77
 - 2.6.5. Additional Resources 83
- 2.7. Kerberos 83
 - 2.7.1. What is Kerberos? 84
 - 2.7.2. Kerberos Terminology 85
 - 2.7.3. How Kerberos Works 87
 - 2.7.4. Kerberos and PAM 88
 - 2.7.5. Configuring a Kerberos 5 Server 88
 - 2.7.6. Configuring a Kerberos 5 Client 90
 - 2.7.7. Domain-to-Realm Mapping 91
 - 2.7.8. Setting Up Secondary KDCs 92
 - 2.7.9. Setting Up Cross Realm Authentication 93
 - 2.7.10. Additional Resources 96
- 2.8. Virtual Private Networks (VPNs) 98
 - 2.8.1. How Does a VPN Work? 98
 - 2.8.2. VPNs and Fedora 99
 - 2.8.3. IPsec 99
 - 2.8.4. Creating an IPsec Connection 99
 - 2.8.5. IPsec Installation 99
 - 2.8.6. IPsec Host-to-Host Configuration 100
 - 2.8.7. IPsec Network-to-Network Configuration 106
 - 2.8.8. Starting and Stopping an IPsec Connection 112
- 2.9. Firewalls 112
 - 2.9.1. Netfilter and IPTables 114
 - 2.9.2. Basic Firewall Configuration 114
 - 2.9.3. Using IPTables 117
 - 2.9.4. Common IPTables Filtering 119
 - 2.9.5. FORWARD and NAT Rules 120
 - 2.9.6. Malicious Software and Spoofed IP Addresses 122
 - 2.9.7. IPTables and Connection Tracking 123
 - 2.9.8. IPv6 123
 - 2.9.9. Additional Resources 124
- 2.10. IPTables 124
 - 2.10.1. Packet Filtering 125
 - 2.10.2. Command Options for IPTables 126
 - 2.10.3. Saving IPTables Rules 135

2.10.4. IPTables Control Scripts	135
2.10.5. IPTables and IPv6	137
2.10.6. Additional Resources	138
3. Encryption	139
3.1. Data at Rest	139
3.2. Full Disk Encryption	139
3.3. File Based Encryption	139
3.4. Data in Motion	139
3.5. Virtual Private Networks	140
3.6. Secure Shell	140
3.7. LUKS Disk Encryption	140
3.7.1. LUKS Implementation in Fedora	140
3.7.2. Manually Encrypting Directories	141
3.7.3. Step-by-Step Instructions	141
3.7.4. What you have just accomplished.	142
3.7.5. Links of Interest	142
3.8. 7-Zip Encrypted Archives	142
3.8.1. 7-Zip Installation in Fedora	142
3.8.2. Step-by-Step Installation Instructions	142
3.8.3. Step-by-Step Usage Instructions	143
3.8.4. Things of note	143
3.9. Using GNU Privacy Guard (GnuPG)	143
3.9.1. Creating GPG Keys in GNOME	144
3.9.2. Creating GPG Keys in KDE	144
3.9.3. Creating GPG Keys Using the Command Line	144
3.9.4. About Public Key Encryption	146
4. General Principles of Information Security	147
4.1. Tips, Guides, and Tools	147
5. Secure Installation	149
5.1. Disk Partitions	149
5.2. Utilize LUKS Partition Encryption	149
6. Software Maintenance	151
6.1. Install Minimal Software	151
6.2. Plan and Configure Security Updates	151
6.3. Adjusting Automatic Updates	151
6.4. Install Signed Packages from Well Known Repositories	151
7. References	153

Preface

1. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the *Liberation Fonts*¹ set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

1.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keycaps and key combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a keycap, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from keycaps by the hyphen connecting each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F1** to switch to the first virtual terminal. Press **Ctrl+Alt+F7** to return to your X-Windows session.

The first paragraph highlights the particular keycap to press. The second highlights two key combinations (each a set of three keycaps with each set pressed simultaneously).

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

¹ <https://fedorahosted.org/liberation-fonts/>

Choose **System > Preferences > Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications > Accessories > Character Map** from the main menu bar. Next, choose **Search > Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit > Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Note the **>** shorthand used to indicate traversal through a menu and its sub-menus. This avoids difficult-to-follow phrasing such as 'Select **Mouse** from the **Preferences** sub-menu in the **System** menu of the main menu bar'.

Mono-spaced Bold Italic or ***Proportional Bold Italic***

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh *john@example.com***.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount */home***.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — *username*, *domain.name*, *file-system*, *package*, *version* and *release*. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

When the Apache HTTP Server accepts requests, it dispatches child processes or threads to handle them. This group of child processes or threads is known as a *server-pool*. Under Apache HTTP Server 2.0, the responsibility for creating and maintaining these server-pools has been abstracted to a group of modules called *Multi-Processing Modules (MPMs)*. Unlike other modules, only one module from the MPM group can be loaded by the Apache HTTP Server.

1.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in mono-spaced roman and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in mono-spaced roman but add syntax highlighting as follows:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo            echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

1.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' won't cause data loss but may cause irritation and frustration.



Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

2. We Need Feedback!

More information about the Linux Security Guide project can be found at <https://fedorahosted.org/securityguide>

To provide feedback for the Security Guide, please file a bug in https://bugzilla.redhat.com/enter_bug.cgi?component=security-guide&product=Fedora%20Documentation. Please select the proper component in the dropdown menu which should be the page name.

Security Overview

Because of the increased reliance on powerful, networked computers to help run businesses and keep track of our personal information, entire industries have been formed around the practice of network and computer security. Enterprises have solicited the knowledge and skills of security experts to properly audit systems and tailor solutions to fit the operating requirements of the organization. Because most organizations are increasingly dynamic in nature, with workers accessing company IT resources locally and remotely, the need for secure computing environments has become more pronounced.

Unfortunately, most organizations (as well as individual users) regard security as an afterthought, a process that is overlooked in favor of increased power, productivity, and budgetary concerns. Proper security implementation is often enacted postmortem — *after* an unauthorized intrusion has already occurred. Security experts agree that taking the correct measures prior to connecting a site to an untrusted network, such as the Internet, is an effective means of thwarting most attempts at intrusion.

1.1. Introduction to Security

1.1.1. What is Computer Security?

Computer security is a general term that covers a wide area of computing and information processing. Industries that depend on computer systems and networks to conduct daily business transactions and access crucial information regard their data as an important part of their overall assets. Several terms and metrics have entered our daily business vocabulary, such as total cost of ownership (TCO) and quality of service (QoS). Using these metrics, industries can calculate aspects such as data integrity and high-availability as part of their planning and process management costs. In some industries, such as electronic commerce, the availability and trustworthiness of data can be the difference between success and failure.

1.1.1.1. How did Computer Security Come about?

Information security has evolved over the years due to the increasing reliance on public networks not to disclose personal, financial, and other restricted information. There are numerous instances such as the Mitnick ¹ and the Vladimir Levin ² cases that prompted organizations across all industries to re-think the way they handle information, as well as its transmission and disclosure. The popularity of the Internet was one of the most important developments that prompted an intensified effort in data security.

An ever-growing number of people are using their personal computers to gain access to the resources that the Internet has to offer. From research and information retrieval to electronic mail and commerce transaction, the Internet has been regarded as one of the most important developments of the 20th century.

The Internet and its earlier protocols, however, were developed as a *trust-based* system. That is, the Internet Protocol was not designed to be secure in itself. There are no approved security standards built into the TCP/IP communications stack, leaving it open to potentially malicious users and processes across the network. Modern developments have made Internet communication more secure, but there are still several incidents that gain national attention and alert us to the fact that nothing is completely safe.

<http://law.jrank.org/pages/3791/Kevin-Mitnick-Case-1999.html>
http://www.livinginternet.com/ii/ia_hackers_levin.htm

1.1.1.2. Security Today

In February of 2000, a Distributed Denial of Service (DDoS) attack was unleashed on several of the most heavily-trafficked sites on the Internet. The attack rendered yahoo.com, cnn.com, amazon.com, fbi.gov, and several other sites completely unreachable to normal users, as it tied up routers for several hours with large-byte ICMP packet transfers, also called a *ping flood*. The attack was brought on by unknown assailants using specially created, widely available programs that scanned vulnerable network servers, installed client applications called *trojans* on the servers, and timed an attack with every infected server flooding the victim sites and rendering them unavailable. Many blame the attack on fundamental flaws in the way routers and the protocols used are structured to accept all incoming data, no matter where or for what purpose the packets are sent.

In 2007, a data breach exploiting the widely-known weaknesses of the Wired Equivalent Privacy (WEP) wireless encryption protocol resulted in the theft from a global financial institution of over 45 million credit card numbers.³

In a separate incident, the billing records of over 2.2 million patients stored on a backup tape were stolen from the front seat of a courier's car.⁴

Currently, an estimated 1.4 billion people use or have used the Internet worldwide.⁵ At the same time:

- On any given day, there are approximately 225 major incidences of security breach reported to the CERT Coordination Center at Carnegie Mellon University.⁶
- In 2003, the number of CERT reported incidences jumped to 137,529 from 82,094 in 2002 and from 52,658 in 2001.⁷
- The worldwide economic impact of the three most dangerous Internet Viruses of the last three years was estimated at US\$13.2 Billion.⁸

From a 2008 global survey of business and technology executives "The Global State of Information Security"⁹, undertaken by *CIO Magazine*, some points are:

- Just 43% of respondents audit or monitor user compliance with security policies
- Only 22% keep an inventory of the outside companies that use their data
- The source of nearly half of security incidents was marked as "Unknown"
- 44% of respondents plan to increase security spending in the next year
- 59% have an information security strategy

These results enforce the reality that computer security has become a quantifiable and justifiable expense for IT budgets. Organizations that require data integrity and high availability elicit the skills of system administrators, developers, and engineers to ensure 24x7 reliability of their systems, services, and information. Falling victim to malicious users, processes, or coordinated attacks is a direct threat to the success of the organization.

Unfortunately, system and network security can be a difficult proposition, requiring an intricate knowledge of how an organization regards, uses, manipulates, and transmits its information.

http://www.theregister.co.uk/2007/05/04/tjx_nonfeasance/

<http://www.healthcareitnews.com/story.cms?id=9408>

<http://www.internetworldstats.com/stats.htm>

http://www.csoonline.com/article/454939/The_Global_State_of_Information_Security_

Understanding the way an organization (and the people that make up the organization) conducts business is paramount to implementing a proper security plan.

1.1.1.3. Standardizing Security

Enterprises in every industry rely on regulations and rules that are set by standards-making bodies such as the American Medical Association (AMA) or the Institute of Electrical and Electronics Engineers (IEEE). The same ideals hold true for information security. Many security consultants and vendors agree upon the standard security model known as CIA, or *Confidentiality, Integrity, and Availability*. This three-tiered model is a generally accepted component to assessing risks of sensitive information and establishing security policy. The following describes the CIA model in further detail:

- **Confidentiality** — Sensitive information must be available only to a set of pre-defined individuals. Unauthorized transmission and usage of information should be restricted. For example, confidentiality of information ensures that a customer's personal or financial information is not obtained by an unauthorized individual for malicious purposes such as identity theft or credit fraud.
- **Integrity** — Information should not be altered in ways that render it incomplete or incorrect. Unauthorized users should be restricted from the ability to modify or destroy sensitive information.
- **Availability** — Information should be accessible to authorized users any time that it is needed. Availability is a warranty that information can be obtained with an agreed-upon frequency and timeliness. This is often measured in terms of percentages and agreed to formally in Service Level Agreements (SLAs) used by network service providers and their enterprise clients.

1.1.2. SELinux

Fedora includes an enhancement to the Linux kernel called SELinux, which implements a Mandatory Access Control (MAC) architecture that provides a fine-grained level of control over files, processes, users and applications in the system. Detailed discussion of SELinux is beyond the scope of this document; however, for more information on SELinux and its use in Fedora, refer to the Fedora SELinux User Guide available at <http://docs.fedoraproject.org/selinux-user-guide/>. For more information on configuring and running services in Fedora that are protected by SELinux, refer to the SELinux Managing Confined Services Guide available at <http://docs.fedoraproject.org/selinux-managing-confined-services-guide>¹⁰. Other available resources for SELinux are listed in [Chapter 7, References](#).

1.1.3. Security Controls

Computer security is often divided into three distinct master categories, commonly referred to as *controls*:

- Physical
- Technical
- Administrative

These three broad categories define the main objectives of proper security implementation. Within these controls are sub-categories that further detail the controls and how to implement them.

¹⁰ <http://docs.fedoraproject.org/selinux-managing-confined-services-guide/>

1.1.3.1. Physical Controls

Physical control is the implementation of security measures in a defined structure used to deter or prevent unauthorized access to sensitive material. Examples of physical controls are:

- Closed-circuit surveillance cameras
- Motion or thermal alarm systems
- Security guards
- Picture IDs
- Locked and dead-bolted steel doors
- Biometrics (includes fingerprint, voice, face, iris, handwriting, and other automated methods used to recognize individuals)

1.1.3.2. Technical Controls

Technical controls use technology as a basis for controlling the access and usage of sensitive data throughout a physical structure and over a network. Technical controls are far-reaching in scope and encompass such technologies as:

- Encryption
- Smart cards
- Network authentication
- Access control lists (ACLs)
- File integrity auditing software

1.1.3.3. Administrative Controls

Administrative controls define the human factors of security. They involve all levels of personnel within an organization and determine which users have access to what resources and information by such means as:

- Training and awareness
- Disaster preparedness and recovery plans
- Personnel recruitment and separation strategies
- Personnel registration and accounting

1.1.4. Conclusion

Now that you have learned about the origins, reasons, and aspects of security, you will find it easier to determine the appropriate course of action with regard to Fedora. It is important to know what factors and conditions make up security in order to plan and implement a proper strategy. With this information in mind, the process can be formalized and the path becomes clearer as you delve deeper into the specifics of the security process.

1.2. Vulnerability Assessment

Given time, resources, and motivation, a cracker can break into nearly any system. At the end of the day, all of the security procedures and technologies currently available cannot guarantee that any systems are completely safe from intrusion. Routers help secure gateways to the Internet. Firewalls help secure the edge of the network. Virtual Private Networks safely pass data in an encrypted stream. Intrusion detection systems warn you of malicious activity. However, the success of each of these technologies is dependent upon a number of variables, including:

- The expertise of the staff responsible for configuring, monitoring, and maintaining the technologies.
- The ability to patch and update services and kernels quickly and efficiently.
- The ability of those responsible to keep constant vigilance over the network.

Given the dynamic state of data systems and technologies, securing corporate resources can be quite complex. Due to this complexity, it is often difficult to find expert resources for all of your systems. While it is possible to have personnel knowledgeable in many areas of information security at a high level, it is difficult to retain staff who are experts in more than a few subject areas. This is mainly because each subject area of information security requires constant attention and focus. Information security does not stand still.

1.2.1. Thinking Like the Enemy

Suppose that you administer an enterprise network. Such networks are commonly comprised of operating systems, applications, servers, network monitors, firewalls, intrusion detection systems, and more. Now imagine trying to keep current with each of these. Given the complexity of today's software and networking environments, exploits and bugs are a certainty. Keeping current with patches and updates for an entire network can prove to be a daunting task in a large organization with heterogeneous systems.

Combine the expertise requirements with the task of keeping current, and it is inevitable that adverse incidents occur, systems are breached, data is corrupted, and service is interrupted.

To augment security technologies and aid in protecting systems, networks, and data, you must think like a cracker and gauge the security of your systems by checking for weaknesses. Preventative vulnerability assessments against your own systems and network resources can reveal potential issues that can be addressed before a cracker exploits it.

A vulnerability assessment is an internal audit of your network and system security; the results of which indicate the confidentiality, integrity, and availability of your network (as explained in [Section 1.1.1.3, "Standardizing Security"](#)). Typically, vulnerability assessment starts with a reconnaissance phase, during which important data regarding the target systems and resources is gathered. This phase leads to the system readiness phase, whereby the target is essentially checked for all known vulnerabilities. The readiness phase culminates in the reporting phase, where the findings are classified into categories of high, medium, and low risk; and methods for improving the security (or mitigating the risk of vulnerability) of the target are discussed.

If you were to perform a vulnerability assessment of your home, you would likely check each door to your home to see if they are closed and locked. You would also check every window, making sure that they closed completely and latch correctly. This same concept applies to systems, networks, and electronic data. Malicious users are the thieves and vandals of your data. Focus on their tools, mentality, and motivations, and you can then react swiftly to their actions.

1.2.2. Defining Assessment and Testing

Vulnerability assessments may be broken down into one of two types: *Outside looking in* and *inside looking around*.

When performing an outside looking in vulnerability assessment, you are attempting to compromise your systems from the outside. Being external to your company provides you with the cracker's viewpoint. You see what a cracker sees — publicly-routable IP addresses, systems on your *DMZ*, external interfaces of your firewall, and more. *DMZ* stands for "demilitarized zone", which corresponds to a computer or small subnetwork that sits between a trusted internal network, such as a corporate private LAN, and an untrusted external network, such as the public Internet. Typically, the *DMZ* contains devices accessible to Internet traffic, such as Web (HTTP) servers, FTP servers, SMTP (e-mail) servers and DNS servers.

When you perform an inside looking around vulnerability assessment, you are somewhat at an advantage since you are internal and your status is elevated to trusted. This is the viewpoint you and your co-workers have once logged on to your systems. You see print servers, file servers, databases, and other resources.

There are striking distinctions between these two types of vulnerability assessments. Being internal to your company gives you elevated privileges more so than any outsider. Still today in most organizations, security is configured in such a manner as to keep intruders out. Very little is done to secure the internals of the organization (such as departmental firewalls, user-level access controls, authentication procedures for internal resources, and more). Typically, there are many more resources when looking around inside as most systems are internal to a company. Once you set yourself outside of the company, you immediately are given an untrusted status. The systems and resources available to you externally are usually very limited.

Consider the difference between vulnerability assessments and *penetration tests*. Think of a vulnerability assessment as the first step to a penetration test. The information gleaned from the assessment is used for testing. Whereas the assessment is undertaken to check for holes and potential vulnerabilities, the penetration testing actually attempts to exploit the findings.

Assessing network infrastructure is a dynamic process. Security, both information and physical, is dynamic. Performing an assessment shows an overview, which can turn up false positives and false negatives.

Security administrators are only as good as the tools they use and the knowledge they retain. Take any of the assessment tools currently available, run them against your system, and it is almost a guarantee that there are some false positives. Whether by program fault or user error, the result is the same. The tool may find vulnerabilities which in reality do not exist (false positive); or, even worse, the tool may not find vulnerabilities that actually do exist (false negative).

Now that the difference between a vulnerability assessment and a penetration test is defined, take the findings of the assessment and review them carefully before conducting a penetration test as part of your new best practices approach.



Warning

Attempting to exploit vulnerabilities on production resources can have adverse effects to the productivity and efficiency of your systems and network.

The following list examines some of the benefits to performing vulnerability assessments.

- Creates proactive focus on information security
- Finds potential exploits before crackers find them
- Results in systems being kept up to date and patched
- Promotes growth and aids in developing staff expertise
- Abates financial loss and negative publicity

1.2.2.1. Establishing a Methodology

To aid in the selection of tools for a vulnerability assessment, it is helpful to establish a vulnerability assessment methodology. Unfortunately, there is no predefined or industry approved methodology at this time; however, common sense and best practices can act as a sufficient guide.

What is the target? Are we looking at one server, or are we looking at our entire network and everything within the network? Are we external or internal to the company? The answers to these questions are important as they help determine not only which tools to select but also the manner in which they are used.

To learn more about establishing methodologies, refer to the following websites:

- <http://www.isecom.org/osstmm/> *The Open Source Security Testing Methodology Manual (OSSTMM)*
- <http://www.owasp.org/> *The Open Web Application Security Project*

1.2.3. Evaluating the Tools

An assessment can start by using some form of an information gathering tool. When assessing the entire network, map the layout first to find the hosts that are running. Once located, examine each host individually. Focusing on these hosts requires another set of tools. Knowing which tools to use may be the most crucial step in finding vulnerabilities.

Just as in any aspect of everyday life, there are many different tools that perform the same job. This concept applies to performing vulnerability assessments as well. There are tools specific to operating systems, applications, and even networks (based on the protocols used). Some tools are free; others are not. Some tools are intuitive and easy to use, while others are cryptic and poorly documented but have features that other tools do not.

Finding the right tools may be a daunting task and in the end, experience counts. If possible, set up a test lab and try out as many tools as you can, noting the strengths and weaknesses of each. Review the README file or man page for the tool. Additionally, look to the Internet for more information, such as articles, step-by-step guides, or even mailing lists specific to a tool.

The tools discussed below are just a small sampling of the available tools.

1.2.3.1. Scanning Hosts with Nmap

Nmap is a popular tool included in Fedora that can be used to determine the layout of a network. Nmap has been available for many years and is probably the most often used tool when gathering information. An excellent man page is included that provides a detailed description of its options and usage. Administrators can use Nmap on a network to find host systems and open ports on those systems.

Nmap is a competent first step in vulnerability assessment. You can map out all the hosts within your network and even pass an option that allows Nmap to attempt to identify the operating system running on a particular host. Nmap is a good foundation for establishing a policy of using secure services and stopping unused services.

1.2.3.1.1. Using Nmap

Nmap can be run from a shell prompt by typing the **nmap** command followed by the hostname or IP address of the machine to scan.

```
nmap foo.example.com
```

The results of a basic scan (which could take up to a few minutes, depending on where the host is located and other network conditions) should look similar to the following:

```
Starting Nmap 4.68 ( http://nmap.org )
Interesting ports on foo.example.com:
Not shown: 1710 filtered ports
PORT      STATE SERVICE
22/tcp    open  ssh
53/tcp    open  domain
70/tcp    closed gopher
80/tcp    open  http
113/tcp   closed auth
```

Nmap tests the most common network communication ports for listening or waiting services. This knowledge can be helpful to an administrator who wants to close down unnecessary or unused services.

For more information about using Nmap, refer to the official homepage at the following URL:

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

1.2.3.2. Nessus

Nessus is a full-service security scanner. The plug-in architecture of Nessus allows users to customize it for their systems and networks. As with any scanner, Nessus is only as good as the signature database it relies upon. Fortunately, Nessus is frequently updated and features full reporting, host scanning, and real-time vulnerability searches. Remember that there could be false positives and false negatives, even in a tool as powerful and as frequently updated as Nessus.



Note

The Nessus client and server software is included in Fedora repositories but requires a subscription to use. It has been included in this document as a reference to users who may be interested in using this popular application.

For more information about Nessus, refer to the official website at the following URL:

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

1.2.3.3. Nikto

Nikto is an excellent common gateway interface (CGI) script scanner. Nikto not only checks for CGI vulnerabilities but does so in an evasive manner, so as to elude intrusion detection systems. It comes with thorough documentation which should be carefully reviewed prior to running the program. If you have Web servers serving up CGI scripts, Nikto can be an excellent resource for checking the security of these servers.

More information about Nikto can be found at the following URL:

<http://www.cirt.net/code/nikto.shtml>

1.2.3.4. VLAD the Scanner

VLAD is a vulnerabilities scanner developed by the RAZOR team at Bindview, Inc., which checks for the SANS Top Ten list of common security issues (SNMP issues, file sharing issues, etc.). While not as full-featured as Nessus, VLAD is worth investigating.



Note

VLAD is not included with Fedora and is not supported. It has been included in this document as a reference to users who may be interested in using this popular application.

More information about VLAD can be found on the RAZOR team website at the following URL:

<http://www.bindview.com/Support/Razor/Utilities/>

1.2.3.5. Anticipating Your Future Needs

Depending upon your target and resources, there are many tools available. There are tools for wireless networks, Novell networks, Windows systems, Linux systems, and more. Another essential part of performing assessments may include reviewing physical security, personnel screening, or voice/PBX network assessment. New concepts, such as *war walking*, which involves scanning the perimeter of your enterprise's physical structures for wireless network vulnerabilities, are some emerging concepts that you can investigate and, if needed, incorporate into your assessments. Imagination and exposure are the only limits of planning and conducting vulnerability assessments.

1.3. Attackers and Vulnerabilities

To plan and implement a good security strategy, first be aware of some of the issues which determined, motivated attackers exploit to compromise systems. However, before detailing these issues, the terminology used when identifying an attacker must be defined.

1.3.1. A Quick History of Hackers

The modern meaning of the term *hacker* has origins dating back to the 1960s and the Massachusetts Institute of Technology (MIT) Tech Model Railroad Club, which designed train sets of large scale and intricate detail. Hacker was a name used for club members who discovered a clever trick or workaround for a problem.

The term hacker has since come to describe everything from computer buffs to gifted programmers. A common trait among most hackers is a willingness to explore in detail how computer systems and

networks function with little or no outside motivation. Open source software developers often consider themselves and their colleagues to be hackers, and use the word as a term of respect.

Typically, hackers follow a form of the *hacker ethic* which dictates that the quest for information and expertise is essential, and that sharing this knowledge is the hackers duty to the community. During this quest for knowledge, some hackers enjoy the academic challenges of circumventing security controls on computer systems. For this reason, the press often uses the term hacker to describe those who illicitly access systems and networks with unscrupulous, malicious, or criminal intent. The more accurate term for this type of computer hacker is *cracker* — a term created by hackers in the mid-1980s to differentiate the two communities.

1.3.1.1. Shades of Gray

Within the community of individuals who find and exploit vulnerabilities in systems and networks are several distinct groups. These groups are often described by the shade of hat that they "wear" when performing their security investigations and this shade is indicative of their intent.

The *white hat hacker* is one who tests networks and systems to examine their performance and determine how vulnerable they are to intrusion. Usually, white hat hackers crack their own systems or the systems of a client who has specifically employed them for the purposes of security auditing. Academic researchers and professional security consultants are two examples of white hat hackers.

A *black hat hacker* is synonymous with a cracker. In general, crackers are less focused on programming and the academic side of breaking into systems. They often rely on available cracking programs and exploit well known vulnerabilities in systems to uncover sensitive information for personal gain or to inflict damage on the target system or network.

The *gray hat hacker*, on the other hand, has the skills and intent of a white hat hacker in most situations but uses his knowledge for less than noble purposes on occasion. A gray hat hacker can be thought of as a white hat hacker who wears a black hat at times to accomplish his own agenda.

Gray hat hackers typically subscribe to another form of the hacker ethic, which says it is acceptable to break into systems as long as the hacker does not commit theft or breach confidentiality. Some would argue, however, that the act of breaking into a system is in itself unethical.

Regardless of the intent of the intruder, it is important to know the weaknesses a cracker may likely attempt to exploit. The remainder of the chapter focuses on these issues.

1.3.2. Threats to Network Security

Bad practices when configuring the following aspects of a network can increase the risk of attack.

1.3.2.1. Insecure Architectures

A misconfigured network is a primary entry point for unauthorized users. Leaving a trust-based, open local network vulnerable to the highly-insecure Internet is much like leaving a door ajar in a crime-ridden neighborhood — nothing may happen for an arbitrary amount of time, but *eventually* someone exploits the opportunity.

1.3.2.1.1. Broadcast Networks

System administrators often fail to realize the importance of networking hardware in their security schemes. Simple hardware such as hubs and routers rely on the broadcast or non-switched principle; that is, whenever a node transmits data across the network to a recipient node, the hub or router

sends a broadcast of the data packets until the recipient node receives and processes the data. This method is the most vulnerable to address resolution protocol (ARP) or media access control (MAC) address spoofing by both outside intruders and unauthorized users on local hosts.

1.3.2.1.2. Centralized Servers

Another potential networking pitfall is the use of centralized computing. A common cost-cutting measure for many businesses is to consolidate all services to a single powerful machine. This can be convenient as it is easier to manage and costs considerably less than multiple-server configurations. However, a centralized server introduces a single point of failure on the network. If the central server is compromised, it may render the network completely useless or worse, prone to data manipulation or theft. In these situations, a central server becomes an open door which allows access to the entire network.

1.3.3. Threats to Server Security

Server security is as important as network security because servers often hold a great deal of an organization's vital information. If a server is compromised, all of its contents may become available for the cracker to steal or manipulate at will. The following sections detail some of the main issues.

1.3.3.1. Unused Services and Open Ports

A full installation of Fedora contains 1000+ application and library packages. However, most server administrators do not opt to install every single package in the distribution, preferring instead to install a base installation of packages, including several server applications.

A common occurrence among system administrators is to install the operating system without paying attention to what programs are actually being installed. This can be problematic because unneeded services may be installed, configured with the default settings, and possibly turned on. This can cause unwanted services, such as Telnet, DHCP, or DNS, to run on a server or workstation without the administrator realizing it, which in turn can cause unwanted traffic to the server, or even, a potential pathway into the system for crackers. Refer To [Section 2.3, "Server Security"](#) for information on closing ports and disabling unused services.

1.3.3.2. Unpatched Services

Most server applications that are included in a default installation are solid, thoroughly tested pieces of software. Having been in use in production environments for many years, their code has been thoroughly refined and many of the bugs have been found and fixed.

However, there is no such thing as perfect software and there is always room for further refinement. Moreover, newer software is often not as rigorously tested as one might expect, because of its recent arrival to production environments or because it may not be as popular as other server software.

Developers and system administrators often find exploitable bugs in server applications and publish the information on bug tracking and security-related websites such as the Bugtraq mailing list (<http://www.securityfocus.com>) or the Computer Emergency Response Team (CERT) website (<http://www.cert.org>). Although these mechanisms are an effective way of alerting the community to security vulnerabilities, it is up to system administrators to patch their systems promptly. This is particularly true because crackers have access to these same vulnerability tracking services and will use the information to crack unpatched systems whenever they can. Good system administration requires vigilance, constant bug tracking, and proper system maintenance to ensure a more secure computing environment.

Refer to [Section 1.5, "Security Updates"](#) for more information about keeping a system up-to-date.

1.3.3.3. Inattentive Administration

Administrators who fail to patch their systems are one of the greatest threats to server security. According to the *SysAdmin, Audit, Network, Security Institute (SANS)*, the primary cause of computer security vulnerability is to "assign untrained people to maintain security and provide neither the training nor the time to make it possible to do the job."¹¹ This applies as much to inexperienced administrators as it does to overconfident or unmotivated administrators.

Some administrators fail to patch their servers and workstations, while others fail to watch log messages from the system kernel or network traffic. Another common error is when default passwords or keys to services are left unchanged. For example, some databases have default administration passwords because the database developers assume that the system administrator changes these passwords immediately after installation. If a database administrator fails to change this password, even an inexperienced cracker can use a widely-known default password to gain administrative privileges to the database. These are only a few examples of how inattentive administration can lead to compromised servers.

1.3.3.4. Inherently Insecure Services

Even the most vigilant organization can fall victim to vulnerabilities if the network services they choose are inherently insecure. For instance, there are many services developed under the assumption that they are used over trusted networks; however, this assumption fails as soon as the service becomes available over the Internet — which is itself inherently untrusted.

One category of insecure network services are those that require unencrypted usernames and passwords for authentication. Telnet and FTP are two such services. If packet sniffing software is monitoring traffic between the remote user and such a service usernames and passwords can be easily intercepted.

Inherently, such services can also more easily fall prey to what the security industry terms the *man-in-the-middle* attack. In this type of attack, a cracker redirects network traffic by tricking a cracked name server on the network to point to his machine instead of the intended server. Once someone opens a remote session to the server, the attacker's machine acts as an invisible conduit, sitting quietly between the remote service and the unsuspecting user capturing information. In this way a cracker can gather administrative passwords and raw data without the server or the user realizing it.

Another category of insecure services include network file systems and information services such as NFS or NIS, which are developed explicitly for LAN usage but are, unfortunately, extended to include WANs (for remote users). NFS does not, by default, have any authentication or security mechanisms configured to prevent a cracker from mounting the NFS share and accessing anything contained therein. NIS, as well, has vital information that must be known by every computer on a network, including passwords and file permissions, within a plain text ASCII or DBM (ASCII-derived) database. A cracker who gains access to this database can then access every user account on a network, including the administrator's account.

By default, Fedora is released with all such services turned off. However, since administrators often find themselves forced to use these services, careful configuration is critical. Refer to [Section 2.3, "Server Security"](#) for more information about setting up services in a safe manner.

<http://www.sans.org/resources/errors.php>

1.3.4. Threats to Workstation and Home PC Security

Workstations and home PCs may not be as prone to attack as networks or servers, but since they often contain sensitive data, such as credit card information, they are targeted by system crackers. Workstations can also be co-opted without the user's knowledge and used by attackers as "slave" machines in coordinated attacks. For these reasons, knowing the vulnerabilities of a workstation can save users the headache of reinstalling the operating system, or worse, recovering from data theft.

1.3.4.1. Bad Passwords

Bad passwords are one of the easiest ways for an attacker to gain access to a system. For more on how to avoid common pitfalls when creating a password, refer to [Section 2.2.3, "Password Security"](#).

1.3.4.2. Vulnerable Client Applications

Although an administrator may have a fully secure and patched server, that does not mean remote users are secure when accessing it. For instance, if the server offers Telnet or FTP services over a public network, an attacker can capture the plain text usernames and passwords as they pass over the network, and then use the account information to access the remote user's workstation.

Even when using secure protocols, such as SSH, a remote user may be vulnerable to certain attacks if they do not keep their client applications updated. For instance, v.1 SSH clients are vulnerable to an X-forwarding attack from malicious SSH servers. Once connected to the server, the attacker can quietly capture any keystrokes and mouse clicks made by the client over the network. This problem was fixed in the v.2 SSH protocol, but it is up to the user to keep track of what applications have such vulnerabilities and update them as necessary.

[Section 2.2, "Workstation Security"](#) discusses in more detail what steps administrators and home users should take to limit the vulnerability of computer workstations.

1.4. Common Exploits and Attacks

[Table 1.1, "Common Exploits"](#) details some of the most common exploits and entry points used by intruders to access organizational network resources. Key to these common exploits are the explanations of how they are performed and how administrators can properly safeguard their network against such attacks.

Exploit	Description	Notes
Null or Default Passwords	Leaving administrative passwords blank or using a default password set by the product vendor. This is most common in hardware such as routers and firewalls, though some services that run on Linux can contain default administrator passwords (though Fedora 12 does not ship with them).	Commonly associated with networking hardware such as routers, firewalls, VPNs, and network attached storage (NAS) appliances. Common in many legacy operating systems, especially those that bundle services (such as UNIX and Windows.) Administrators sometimes create privileged user accounts in a rush and leave the password null, creating a perfect entry point for malicious users who discover the account.

Exploit	Description	Notes
Default Shared Keys	Secure services sometimes package default security keys for development or evaluation testing purposes. If these keys are left unchanged and are placed in a production environment on the Internet, <i>all</i> users with the same default keys have access to that shared-key resource, and any sensitive information that it contains.	Most common in wireless access points and preconfigured secure server appliances.
IP Spoofing	A remote machine acts as a node on your local network, finds vulnerabilities with your servers, and installs a backdoor program or trojan horse to gain control over your network resources.	Spoofing is quite difficult as it involves the attacker predicting TCP/IP sequence numbers to coordinate a connection to target systems, but several tools are available to assist crackers in performing such a vulnerability. Depends on target system running services (such as rsh , telnet , FTP and others) that use <i>source-based</i> authentication techniques, which are not recommended when compared to PKI or other forms of encrypted authentication used in ssh or SSL/TLS .
Eavesdropping	Collecting data that passes between two active nodes on a network by eavesdropping on the connection between the two nodes.	This type of attack works mostly with plain text transmission protocols such as Telnet, FTP, and HTTP transfers. Remote attacker must have access to a compromised system on a LAN in order to perform such an attack; usually the cracker has used an active attack (such as IP spoofing or man-in-the-middle) to compromise a system on the LAN. Preventative measures include services with cryptographic key exchange, one-time passwords, or encrypted authentication to prevent password snooping; strong encryption during transmission is also advised.
Service Vulnerabilities	An attacker finds a flaw or loophole in a service run over the Internet; through this vulnerability, the attacker compromises the entire system and any data that it may hold, and could possibly compromise other systems on the network.	HTTP-based services such as CGI are vulnerable to remote command execution and even interactive shell access. Even if the HTTP service runs as a non-privileged user such as "nobody", information such as configuration files and network maps can be read, or the attacker can start a denial of service attack which

Exploit	Description	Notes
		<p>drains system resources or renders it unavailable to other users.</p> <p>Services sometimes can have vulnerabilities that go unnoticed during development and testing; these vulnerabilities (such as <i>buffer overflows</i>, where attackers crash a service using arbitrary values that fill the memory buffer of an application, giving the attacker an interactive command prompt from which they may execute arbitrary commands) can give complete administrative control to an attacker.</p> <p>Administrators should make sure that services do not run as the root user, and should stay vigilant of patches and errata updates for applications from vendors or security organizations such as CERT and CVE.</p>
Application Vulnerabilities	<p>Attackers find faults in desktop and workstation applications (such as e-mail clients) and execute arbitrary code, implant trojan horses for future compromise, or crash systems. Further exploitation can occur if the compromised workstation has administrative privileges on the rest of the network.</p>	<p>Workstations and desktops are more prone to exploitation as workers do not have the expertise or experience to prevent or detect a compromise; it is imperative to inform individuals of the risks they are taking when they install unauthorized software or open unsolicited email attachments. Safeguards can be implemented such that email client software does not automatically open or execute attachments. Additionally, the automatic update of workstation software via Red Hat Network or other system management services can alleviate the burdens of multi-seat security deployments.</p>
Denial of Service (DoS) Attacks	<p>Attacker or group of attackers coordinate against an organization's network or server resources by sending unauthorized packets to the target host (either server, router, or workstation). This forces the resource to become unavailable to legitimate users.</p>	<p>The most reported DoS case in the US occurred in 2000. Several highly-trafficked commercial and government sites were rendered unavailable by a coordinated ping flood attack using several compromised systems with high bandwidth connections acting as <i>zombies</i>, or redirected broadcast nodes.</p> <p>Source packets are usually forged (as well as rebroadcasted), making</p>

Exploit	Description	Notes
		investigation as to the true source of the attack difficult. Advances in ingress filtering (IETF rfc2267) using iptables and Network Intrusion Detection Systems such as snort assist administrators in tracking down and preventing distributed DoS attacks.

Table 1.1. Common Exploits


1.5. Security Updates

As security vulnerabilities are discovered, the affected software must be updated in order to limit any potential security risks. If the software is part of a package within a Fedora distribution that is currently supported, Fedora is committed to releasing updated packages that fix the vulnerability as soon as is possible. Often, announcements about a given security exploit are accompanied with a patch (or source code that fixes the problem). This patch is then applied to the Fedora package and tested and released as an errata update. However, if an announcement does not include a patch, a developer first works with the maintainer of the software to fix the problem. Once the problem is fixed, the package is tested and released as an errata update.

If an errata update is released for software used on your system, it is highly recommended that you update the affected packages as soon as possible to minimize the amount of time the system is potentially vulnerable.

1.5.1. Updating Packages

When updating software on a system, it is important to download the update from a trusted source. An attacker can easily rebuild a package with the same version number as the one that is supposed to fix the problem but with a different security exploit and release it on the Internet. If this happens, using security measures such as verifying files against the original RPM does not detect the exploit. Thus, it is very important to only download RPMs from trusted sources, such as from Fedora and to check the signature of the package to verify its integrity.



Note
Fedora includes a convenient panel icon that displays visible alerts when there is an update for a Fedora system.

1.5.2. Verifying Signed Packages

All Fedora packages are signed with the Fedora *GPG* key. GPG stands for GNU Privacy Guard, or GnuPG, a free software package used for ensuring the authenticity of distributed files. For example, a private key (secret key) locks the package while the public key unlocks and verifies the package. If the public key distributed by Fedora does not match the private key during RPM verification, the package may have been altered and therefore cannot be trusted.

The RPM utility within Fedora automatically tries to verify the GPG signature of an RPM package before installing it. If the Fedora GPG key is not installed, install it from a secure, static location, such as an Fedora installation CD-ROM or DVD.

Assuming the disc is mounted in `/mnt/cdrom`, use the following command to import it into the *keyring* (a database of trusted keys on the system):

```
rpm --import /mnt/cdrom/RPM-GPG-KEY
```

To display a list of all keys installed for RPM verification, execute the following command:

```
rpm -qa gpg-pubkey*
```

The output will look similar to the following:

```
gpg-pubkey-db42a60e-37ea5438
```

To display details about a specific key, use the `rpm -qi` command followed by the output from the previous command, as in this example:

```
rpm -qi gpg-pubkey-db42a60e-37ea5438
```

It is extremely important to verify the signature of the RPM files before installing them to ensure that they have not been altered from the original source of the packages. To verify all the downloaded packages at once, issue the following command:

```
rpm -K /tmp/updates/*.rpm
```

For each package, if the GPG key verifies successfully, the command returns `gpg OK`. If it doesn't, make sure you are using the correct Fedora public key, as well as verifying the source of the content. Packages that do not pass GPG verifications should not be installed, as they may have been altered by a third party.

After verifying the GPG key and downloading all the packages associated with the errata report, install the packages as root at a shell prompt.

1.5.3. Installing Signed Packages

Installation for most packages can be done safely (except kernel packages) by issuing the following command:

```
rpm -Uvh /tmp/updates/*.rpm
```

For kernel packages use the following command:

```
rpm -ivh /tmp/updates/<kernel-package>
```

Replace `<kernel-package>` in the previous example with the name of the kernel RPM.

Once the machine has been safely rebooted using the new kernel, the old kernel may be removed using the following command:

```
rpm -e <old-kernel-package>
```

Replace `<old-kernel-package>` in the previous example with the name of the older kernel RPM.



Note

It is not a requirement that the old kernel be removed. The default boot loader, GRUB, allows for multiple kernels to be installed, then chosen from a menu at boot time.



Important

Before installing any security errata, be sure to read any special instructions contained in the errata report and execute them accordingly. Refer to *Section 1.5.4, “Applying the Changes”* for general instructions about applying the changes made by an errata update.

1.5.4. Applying the Changes

After downloading and installing security errata and updates, it is important to halt usage of the older software and begin using the new software. How this is done depends on the type of software that has been updated. The following list itemizes the general categories of software and provides instructions for using the updated versions after a package upgrade.



Note

In general, rebooting the system is the surest way to ensure that the latest version of a software package is used; however, this option is not always required, or available to the system administrator.

Applications

User-space applications are any programs that can be initiated by a system user. Typically, such applications are used only when a user, script, or automated task utility launches them and they do not persist for long periods of time.

Once such a user-space application is updated, halt any instances of the application on the system and launch the program again to use the updated version.

Kernel

The kernel is the core software component for the Fedora operating system. It manages access to memory, the processor, and peripherals as well as schedules all tasks.

Because of its central role, the kernel cannot be restarted without also stopping the computer. Therefore, an updated version of the kernel cannot be used until the system is rebooted.

Shared Libraries

Shared libraries are units of code, such as **glibc**, which are used by a number of applications and services. Applications utilizing a shared library typically load the shared code when the application is initialized, so any applications using the updated library must be halted and relaunched.

To determine which running applications link against a particular library, use the **lssof** command as in the following example:

```
lssof /lib/libwrap.so*
```

This command returns a list of all the running programs which use TCP wrappers for host access control. Therefore, any program listed must be halted and relaunched if the **tcp_wrappers** package is updated.

SysV Services

SysV services are persistent server programs launched during the boot process. Examples of SysV services include **sshd**, **vsftpd**, and **xinetd**.

Because these programs usually persist in memory as long as the machine is booted, each updated SysV service must be halted and relaunched after the package is upgraded. This can be done using the **Services Configuration Tool** or by logging into a root shell prompt and issuing the **/sbin/service** command as in the following example:

```
/sbin/service <service-name> restart
```

In the previous example, replace *<service-name>* with the name of the service, such as **sshd**.

xinetd Services

Services controlled by the **xinetd** super service only run when there is an active connection. Examples of services controlled by **xinetd** include Telnet, IMAP, and POP3.

Because new instances of these services are launched by **xinetd** each time a new request is received, connections that occur after an upgrade are handled by the updated software. However, if there are active connections at the time the **xinetd** controlled service is upgraded, they are serviced by the older version of the software.

To kill off older instances of a particular **xinetd** controlled service, upgrade the package for the service then halt all processes currently running. To determine if the process is running, use the **ps** command and then use the **kill** or **killall** command to halt current instances of the service.

For example, if security errata **imap** packages are released, upgrade the packages, then type the following command as root into a shell prompt:

```
ps -aux | grep imap
```

This command returns all active IMAP sessions. Individual sessions can then be terminated by issuing the following command:

```
kill <PID>
```

If this fails to terminate the session, use the following command instead:

```
kill -9 <PID>
```

In the previous examples, replace *<PID>* with the process identification number (found in the second column of the **ps** command) for an IMAP session.

To kill all active IMAP sessions, issue the following command:

```
killall imapd
```


Securing Your Network

2.1. Local users may install trusted packages



Non-privileged users may install software.

In Fedora 12, a *local* user may install *signed* packages without authentication. This is a change from Fedora 11.

In common use cases, local desktop users frequently install packages. In Fedora 11, this required authentication. In Fedora 11, if the user wishes to install an unsigned package, a second authentication is required. Since the desktop user is typically the owner and sole user of the machine, the default was changed in Fedora 12 to allow a local user to install signed (trusted) packages without authentication. Unsigned packages continue to require authentication.

This change only affects installs and updates made through the graphical interface. It does not affect *yum*, nor does it allow packages to be removed without authentication.

Some administrators may prefer the old behavior. To restore the Fedora 11 behavior, create a file in `/var/lib/polkit-1/localauthority/20-org.d` (name it anything you want as long as it ends with `.pkla`) and the content should be

```
[NoUserSignedInstall]
Identity=unix-user:*
Action=org.freedesktop.packagekit.package-install
ResultAny=no
ResultInactive=no
ResultActive=auth_admin
```

It is important to note that, as of this writing, there is some discussion as to whether this feature may be reverted. There is also a question about whether the above fix works for all users. This document will be updated as new information becomes available.

Those that want to follow the detailed discussion can refer to https://bugzilla.redhat.com/show_bug.cgi?id=534047. Be advised that most of those commenting are developers and frequently have software and understanding beyond ordinary users.

2.2. Workstation Security

Securing a Linux environment begins with the workstation. Whether locking down a personal machine or securing an enterprise system, sound security policy begins with the individual computer. A computer network is only as secure as its weakest node.

2.2.1. Evaluating Workstation Security

When evaluating the security of a Fedora workstation, consider the following:

- *BIOS and Boot Loader Security* — Can an unauthorized user physically access the machine and boot into single user or rescue mode without a password?

- *Password Security* — How secure are the user account passwords on the machine?
- *Administrative Controls* — Who has an account on the system and how much administrative control do they have?
- *Available Network Services* — What services are listening for requests from the network and should they be running at all?
- *Personal Firewalls* — What type of firewall, if any, is necessary?
- *Security Enhanced Communication Tools* — Which tools should be used to communicate between workstations and which should be avoided?

2.2.2. BIOS and Boot Loader Security

Password protection for the BIOS (or BIOS equivalent) and the boot loader can prevent unauthorized users who have physical access to systems from booting using removable media or obtaining root privileges through single user mode. The security measures you should take to protect against such attacks depends both on the sensitivity of the information on the workstation and the location of the machine.

For example, if a machine is used in a trade show and contains no sensitive information, then it may not be critical to prevent such attacks. However, if an employee's laptop with private, unencrypted SSH keys for the corporate network is left unattended at that same trade show, it could lead to a major security breach with ramifications for the entire company.

If the workstation is located in a place where only authorized or trusted people have access, however, then securing the BIOS or the boot loader may not be necessary.

2.2.2.1. BIOS Passwords

The two primary reasons for password protecting the BIOS of a computer are¹:

1. *Preventing Changes to BIOS Settings* — If an intruder has access to the BIOS, they can set it to boot from a diskette or CD-ROM. This makes it possible for them to enter rescue mode or single user mode, which in turn allows them to start arbitrary processes on the system or copy sensitive data.
2. *Preventing System Booting* — Some BIOSes allow password protection of the boot process. When activated, an attacker is forced to enter a password before the BIOS launches the boot loader.

Because the methods for setting a BIOS password vary between computer manufacturers, consult the computer's manual for specific instructions.

If you forget the BIOS password, it can either be reset with jumpers on the motherboard or by disconnecting the CMOS battery. For this reason, it is good practice to lock the computer case if possible. However, consult the manual for the computer or motherboard before attempting to disconnect the CMOS battery.

¹Since system BIOSes differ between manufacturers, some may not support password protection of either type, while others may support one type but not the other.

2.2.2.1.1. Securing Non-x86 Platforms

Other architectures use different programs to perform low-level tasks roughly equivalent to those of the BIOS on x86 systems. For instance, Intel® Itanium™ computers use the *Extensible Firmware Interface (EFI)* shell.

For instructions on password protecting BIOS-like programs on other architectures, refer to the manufacturer's instructions.

2.2.2.2. Boot Loader Passwords

The primary reasons for password protecting a Linux boot loader are as follows:

1. *Preventing Access to Single User Mode* — If attackers can boot the system into single user mode, they are logged in automatically as root without being prompted for the root password.
2. *Preventing Access to the GRUB Console* — If the machine uses GRUB as its boot loader, an attacker can use the GRUB editor interface to change its configuration or to gather information using the **cat** command.
3. *Preventing Access to Insecure Operating Systems* — If it is a dual-boot system, an attacker can select an operating system at boot time (for example, DOS), which ignores access controls and file permissions.

Fedora ships with the GRUB boot loader on the x86 platform. For a detailed look at GRUB, refer to the Red Hat Installation Guide.

2.2.2.2.1. Password Protecting GRUB

You can configure GRUB to address the first two issues listed in [Section 2.2.2.2, “Boot Loader Passwords”](#) by adding a password directive to its configuration file. To do this, first choose a strong password, open a shell, log in as root, and then type the following command:

```
/sbin/grub-md5-crypt
```

When prompted, type the GRUB password and press **Enter**. This returns an MD5 hash of the password.

Next, edit the GRUB configuration file **/boot/grub/grub.conf**. Open the file and below the **timeout** line in the main section of the document, add the following line:

```
password --md5 <password-hash>
```

Replace *<password-hash>* with the value returned by **/sbin/grub-md5-crypt**².

The next time the system boots, the GRUB menu prevents access to the editor or command interface without first pressing **p** followed by the GRUB password.

Unfortunately, this solution does not prevent an attacker from booting into an insecure operating system in a dual-boot environment. For this, a different part of the **/boot/grub/grub.conf** file must be edited.

²GRUB also accepts unencrypted passwords, but it is recommended that an MD5 hash be used for added security.

Look for the `title` line of the operating system that you want to secure, and add a line with the **lock** directive immediately beneath it.

For a DOS system, the stanza should begin similar to the following:

```
title DOS lock
```



Warning

A password line must be present in the main section of the `/boot/grub/grub.conf` file for this method to work properly. Otherwise, an attacker can access the GRUB editor interface and remove the lock line.

To create a different password for a particular kernel or operating system, add a **lock** line to the stanza, followed by a password line.

Each stanza protected with a unique password should begin with lines similar to the following example:

```
title DOS lock password --md5 <password-hash>
```

2.2.3. Password Security

Passwords are the primary method that Fedora uses to verify a user's identity. This is why password security is so important for protection of the user, the workstation, and the network.

For security purposes, the installation program configures the system to use *Message-Digest Algorithm (MD5)* and shadow passwords. It is highly recommended that you do not alter these settings.

If MD5 passwords are deselected during installation, the older *Data Encryption Standard (DES)* format is used. This format limits passwords to eight alphanumeric characters (disallowing punctuation and other special characters), and provides a modest 56-bit level of encryption.

If shadow passwords are deselected during installation, all passwords are stored as a one-way hash in the world-readable `/etc/passwd` file, which makes the system vulnerable to offline password cracking attacks. If an intruder can gain access to the machine as a regular user, he can copy the `/etc/passwd` file to his own machine and run any number of password cracking programs against it. If there is an insecure password in the file, it is only a matter of time before the password cracker discovers it.

Shadow passwords eliminate this type of attack by storing the password hashes in the file `/etc/shadow`, which is readable only by the root user.

This forces a potential attacker to attempt password cracking remotely by logging into a network service on the machine, such as SSH or FTP. This sort of brute-force attack is much slower and leaves an obvious trail as hundreds of failed login attempts are written to system files. Of course, if the cracker starts an attack in the middle of the night on a system with weak passwords, the cracker may have gained access before dawn and edited the log files to cover his tracks.

In addition to format and storage considerations is the issue of content. The single most important thing a user can do to protect his account against a password cracking attack is create a strong password.

2.2.3.1. Creating Strong Passwords

When creating a secure password, it is a good idea to follow these guidelines:

- *Do Not Use Only Words or Numbers* — Never use only numbers or words in a password.

Some insecure examples include the following:

- 8675309
 - juan
 - hackme
- *Do Not Use Recognizable Words* — Words such as proper names, dictionary words, or even terms from television shows or novels should be avoided, even if they are bookended with numbers.

Some insecure examples include the following:

- john1
 - DS-9
 - mentat123
- *Do Not Use Words in Foreign Languages* — Password cracking programs often check against word lists that encompass dictionaries of many languages. Relying on foreign languages for secure passwords is not secure.

Some insecure examples include the following:

- cheguevara
 - bienvenido1
 - 1dumbKopf
- *Do Not Use Hacker Terminology* — If you think you are elite because you use hacker terminology — also called l337 (LEET) speak — in your password, think again. Many word lists include LEET speak.

Some insecure examples include the following:

- H4X0R
 - 1337
- *Do Not Use Personal Information* — Avoid using any personal information in your passwords. If the attacker knows your identity, the task of deducing your password becomes easier. The following is a list of the types of information to avoid when creating a password:

Some insecure examples include the following:

- Your name
- The names of pets

- The names of family members
- Any birth dates
- Your phone number or zip code
- *Do Not Invert Recognizable Words* — Good password checkers always reverse common words, so inverting a bad password does not make it any more secure.

Some insecure examples include the following:

- R0X4H
- nauj
- 9-DS
- *Do Not Write Down Your Password* — Never store a password on paper. It is much safer to memorize it.
- *Do Not Use the Same Password For All Machines* — It is important to make separate passwords for each machine. This way if one system is compromised, all of your machines are not immediately at risk.

The following guidelines will help you to create a strong password:

- *Make the Password at Least Eight Characters Long* — The longer the password, the better. If using MD5 passwords, it should be 15 characters or longer. With DES passwords, use the maximum length (eight characters).
- *Mix Upper and Lower Case Letters* — Fedora is case sensitive, so mix cases to enhance the strength of the password.
- *Mix Letters and Numbers* — Adding numbers to passwords, especially when added to the middle (not just at the beginning or the end), can enhance password strength.
- *Include Non-Alphanumeric Characters* — Special characters such as &, \$, and > can greatly improve the strength of a password (this is not possible if using DES passwords).
- *Pick a Password You Can Remember* — The best password in the world does little good if you cannot remember it; use acronyms or other mnemonic devices to aid in memorizing passwords.

With all these rules, it may seem difficult to create a password that meets all of the criteria for good passwords while avoiding the traits of a bad one. Fortunately, there are some steps you can take to generate an easily-remembered, secure password.

2.2.3.1.1. Secure Password Creation Methodology

There are many methods that people use to create secure passwords. One of the more popular methods involves acronyms. For example:

- Think of an easily-remembered phrase, such as:
"over the river and through the woods, to grandmother's house we go."

- Next, turn it into an acronym (including the punctuation).

otrattw, tghwg.

- Add complexity by substituting numbers and symbols for letters in the acronym. For example, substitute **7** for **t** and the at symbol (**@**) for **a**:

o7r@77w, 7ghwg.

- Add more complexity by capitalizing at least one letter, such as **H**.

o7r@77w, 7gHwg.

- *Finally, do not use the example password above for any systems, ever.*

While creating secure passwords is imperative, managing them properly is also important, especially for system administrators within larger organizations. The following section details good practices for creating and managing user passwords within an organization.

2.2.3.2. Creating User Passwords Within an Organization

If an organization has a large number of users, the system administrators have two basic options available to force the use of good passwords. They can create passwords for the user, or they can let users create their own passwords, while verifying the passwords are of acceptable quality.

Creating the passwords for the users ensures that the passwords are good, but it becomes a daunting task as the organization grows. It also increases the risk of users writing their passwords down.

For these reasons, most system administrators prefer to have the users create their own passwords, but actively verify that the passwords are good and, in some cases, force users to change their passwords periodically through password aging.

2.2.3.2.1. Forcing Strong Passwords

To protect the network from intrusion it is a good idea for system administrators to verify that the passwords used within an organization are strong ones. When users are asked to create or change passwords, they can use the command line application **passwd**, which is *Pluggable Authentication Manager (PAM)* aware and therefore checks to see if the password is too short or otherwise easy to crack. This check is performed using the **pam_cracklib.so** PAM module. Since PAM is customizable, it is possible to add more password integrity checkers, such as **pam_passwdqc** (available from <http://www.openwall.com/passwdqc/>) or to write a new module. For a list of available PAM modules, refer to <http://www.kernel.org/pub/linux/libs/pam/modules.html>. For more information about PAM, refer to [Section 2.5, “Pluggable Authentication Modules \(PAM\)”](#).

The password check that is performed at the time of their creation does not discover bad passwords as effectively as running a password cracking program against the passwords.

Many password cracking programs are available that run under Fedora, although none ship with the operating system. Below is a brief list of some of the more popular password cracking programs:

- **John The Ripper** — A fast and flexible password cracking program. It allows the use of multiple word lists and is capable of brute-force password cracking. It is available online at <http://www.openwall.com/john/>.

- **Crack** — Perhaps the most well known password cracking software, **Crack** is also very fast, though not as easy to use as **John The Ripper**. It can be found online at <http://www.crypticide.com/alecm/security/crack/c50-faq.html>.
- **Slurpie** — **Slurpie** is similar to **John The Ripper** and **Crack**, but it is designed to run on multiple computers simultaneously, creating a distributed password cracking attack. It can be found along with a number of other distributed attack security evaluation tools online at <http://www.ussrback.com/distributed.htm>.



Warning

Always get authorization in writing before attempting to crack passwords within an organization.

2.2.3.2.2. Passphrases

Passphrases and passwords are the cornerstone to security in most of today's systems. Unfortunately, techniques such as biometrics and two-factor authentication have not yet become mainstream in many systems. If passwords are going to be used to secure a system, then the use of passphrases should be considered. Passphrases are longer than passwords and provide better protection than a password even when implemented with non-standard characters such as numbers and symbols.

2.2.3.2.3. Password Aging

Password aging is another technique used by system administrators to defend against bad passwords within an organization. Password aging means that after a specified period (usually 90 days), the user is prompted to create a new password. The theory behind this is that if a user is forced to change his password periodically, a cracked password is only useful to an intruder for a limited amount of time. The downside to password aging, however, is that users are more likely to write their passwords down.

There are two primary programs used to specify password aging under Fedora: the **chage** command or the graphical **User Manager (system-config-users)** application.

The **-M** option of the **chage** command specifies the maximum number of days the password is valid. For example, to set a user's password to expire in 90 days, use the following command:

```
chage -M 90 <username>
```

In the above command, replace **<username>** with the name of the user. To disable password expiration, it is traditional to use a value of **99999** after the **-M** option (this equates to a little over 273 years).

You can also use the **chage** command in interactive mode to modify multiple password aging and account details. Use the following command to enter interactive mode:

```
chage <username>
```

The following is a sample interactive session using this command:

```
[root@myServer ~]# chage davido
Changing the aging information for davido
Enter the new value, or press ENTER for the default
Minimum Password Age [0]: 10
```

```

Maximum Password Age [99999]: 90
Last Password Change (YYYY-MM-DD) [2006-08-18]:
Password Expiration Warning [7]:
Password Inactive [-1]:
Account Expiration Date (YYYY-MM-DD) [1969-12-31]:
[root@myServer ~]#

```

Refer to the man page for `chage` for more information on the available options.

You can also use the graphical **User Manager** application to create password aging policies, as follows. Note: you need Administrator privileges to perform this procedure.

1. Click the **System** menu on the Panel, point to **Administration** and then click **Users and Groups** to display the User Manager. Alternatively, type the command **system-config-users** at a shell prompt.
2. Click the **Users** tab, and select the required user in the list of users.
3. Click **Properties** on the toolbar to display the User Properties dialog box (or choose **Properties** on the **File** menu).
4. Click the **Password Info** tab, and select the check box for **Enable password expiration**.
5. Enter the required value in the **Days before change required** field, and click **OK**.

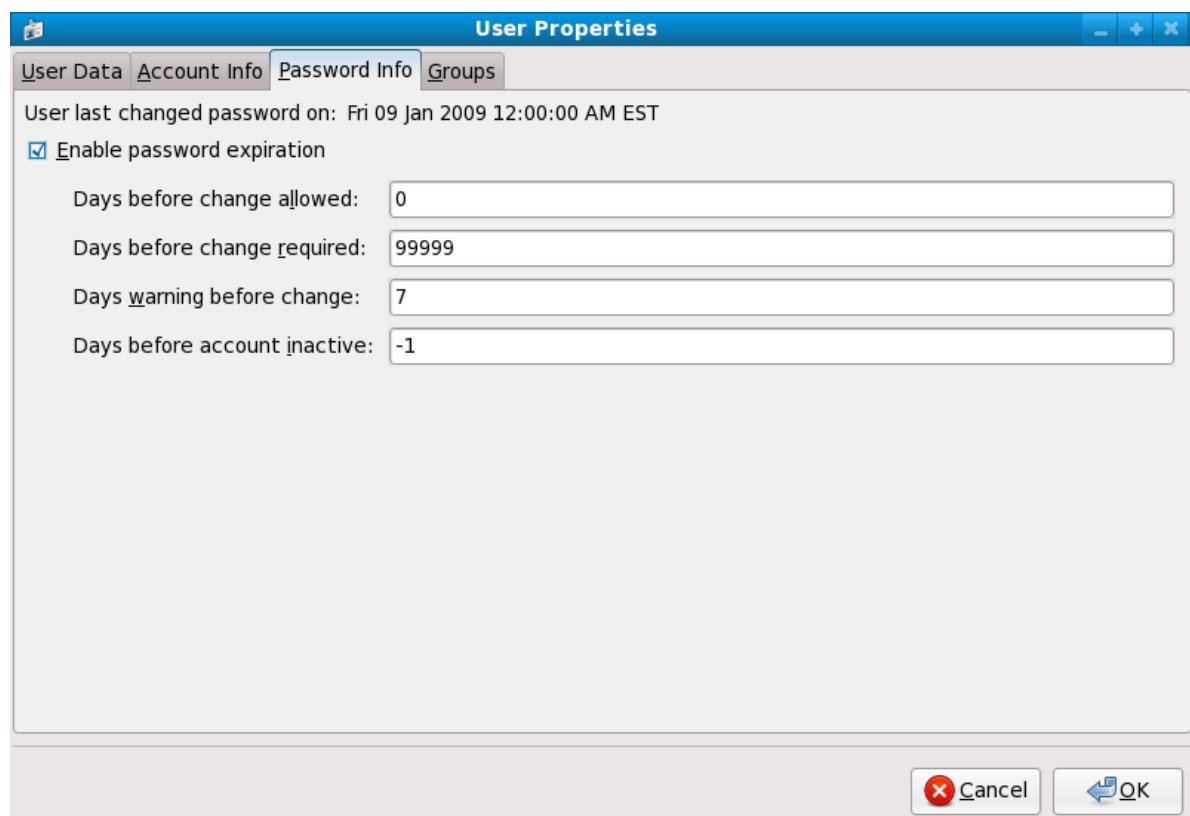


Figure 2.1. Specifying password aging options

2.2.4. Administrative Controls

When administering a home machine, the user must perform some tasks as the root user or by acquiring effective root privileges via a *setuid* program, such as **sudo** or **su**. A *setuid* program is

one that operates with the user ID (*UID*) of the program's owner rather than the user operating the program. Such programs are denoted by an *s* in the owner section of a long format listing, as in the following example:

```
-rwsr-xr-x 1 root root 47324 May 1 08:09 /bin/su
```



Note

The *s* may be upper case or lower case. If it appears as upper case, it means that the underlying permission bit has not been set.

For the system administrators of an organization, however, choices must be made as to how much administrative access users within the organization should have to their machine. Through a PAM module called **pam_console.so**, some activities normally reserved only for the root user, such as rebooting and mounting removable media are allowed for the first user that logs in at the physical console (refer to [Section 2.5, “Pluggable Authentication Modules \(PAM\)”](#) for more information about the **pam_console.so** module.) However, other important system administration tasks, such as altering network settings, configuring a new mouse, or mounting network devices, are not possible without administrative privileges. As a result, system administrators must decide how much access the users on their network should receive.

2.2.4.1. Allowing Root Access

If the users within an organization are trusted and computer-literate, then allowing them root access may not be an issue. Allowing root access by users means that minor activities, like adding devices or configuring network interfaces, can be handled by the individual users, leaving system administrators free to deal with network security and other important issues.

On the other hand, giving root access to individual users can lead to the following issues:

- *Machine Misconfiguration* — Users with root access can misconfigure their machines and require assistance to resolve issues. Even worse, they might open up security holes without knowing it.
- *Running Insecure Services* — Users with root access might run insecure servers on their machine, such as FTP or Telnet, potentially putting usernames and passwords at risk. These services transmit this information over the network in plain text.
- *Running Email Attachments As Root* — Although rare, email viruses that affect Linux do exist. The only time they are a threat, however, is when they are run by the root user.

2.2.4.2. Disallowing Root Access

If an administrator is uncomfortable allowing users to log in as root for these or other reasons, the root password should be kept secret, and access to runlevel one or single user mode should be disallowed through boot loader password protection (refer to [Section 2.2.2.2, “Boot Loader Passwords”](#) for more information on this topic.)

[Table 2.1, “Methods of Disabling the Root Account”](#) describes ways that an administrator can further ensure that root logins are disallowed:

Method	Description	Effects	Does Not Affect
Changing the root shell.	Edit the /etc/passwd file and change the shell from /bin/bash to /sbin/nologin .	Prevents access to the root shell and logs any such attempts. The following programs are prevented from accessing the root account: <ul style="list-style-type: none"> · login · gdm · kdm · xdm · su · ssh · scp · sftp 	Programs that do not require a shell, such as FTP clients, mail clients, and many setuid programs. The following programs are <i>not</i> prevented from accessing the root account: <ul style="list-style-type: none"> · sudo · FTP clients · Email clients
Disabling root access via any console device (tty).	An empty /etc/securetty file prevents root login on any devices attached to the computer.	Prevents access to the root account via the console or the network. The following programs are prevented from accessing the root account: <ul style="list-style-type: none"> · login · gdm · kdm · xdm · Other network services that open a tty 	Programs that do not log in as root, but perform administrative tasks through setuid or other mechanisms. The following programs are <i>not</i> prevented from accessing the root account: <ul style="list-style-type: none"> · su · sudo · ssh · scp · sftp
Disabling root SSH logins.	Edit the /etc/ssh/sshd_config file and set the PermitRootLogin parameter to no .	Prevents root access via the OpenSSH suite of tools. The following programs are prevented from accessing the root account: <ul style="list-style-type: none"> · ssh · scp · sftp 	This only prevents root access to the OpenSSH suite of tools.
Use PAM to limit root access to services.	Edit the file for the target service in the /etc/pam.d/ directory. Make sure the pam_listfile.so is required for authentication. ¹	Prevents root access to network services that are PAM aware. The following services are prevented from accessing the root account: <ul style="list-style-type: none"> · FTP clients · Email clients · login · gdm · kdm · xdm · ssh · scp 	Programs and services that are not PAM aware.


Method	Description	Effects	Does Not Affect
		<ul style="list-style-type: none"> · sftp · Any PAM aware services 	

Refer to [Section 2.2.4.2.4, “Disabling Root Using PAM”](#) for details.

Table 2.1. Methods of Disabling the Root Account

2.2.4.2.1. Disabling the Root Shell

To prevent users from logging in directly as root, the system administrator can set the root account's shell to **/sbin/nologin** in the **/etc/passwd** file. This prevents access to the root account through commands that require a shell, such as the **su** and the **ssh** commands.




Important
 Programs that do not require access to the shell, such as email clients or the **sudo** command, can still access the root account.

2.2.4.2.2. Disabling Root Logins

To further limit access to the root account, administrators can disable root logins at the console by editing the **/etc/securetty** file. This file lists all devices the root user is allowed to log into. If the file does not exist at all, the root user can log in through any communication device on the system, whether via the console or a raw network interface. This is dangerous, because a user can log in to his machine as root via Telnet, which transmits the password in plain text over the network. By default, Fedora's **/etc/securetty** file only allows the root user to log in at the console physically attached to the machine. To prevent root from logging in, remove the contents of this file by typing the following command:

```
echo > /etc/securetty
```



Warning
 A blank **/etc/securetty** file does *not* prevent the root user from logging in remotely using the OpenSSH suite of tools because the console is not opened until after authentication.

2.2.4.2.3. Disabling Root SSH Logins

Root logins via the SSH protocol are disabled by default in Fedora; however, if this option has been enabled, it can be disabled again by editing the SSH daemon's configuration file (**/etc/ssh/sshd_config**). Change the line that reads:

```
PermitRootLogin yes
```

to read as follows:

```
PermitRootLogin no
```

For these changes to take effect, the SSH daemon must be restarted. This can be done via the following command:

```
kill -HUP `cat /var/run/sshd.pid`
```

2.2.4.2.4. Disabling Root Using PAM

PAM, through the `/lib/security/pam_listfile.so` module, allows great flexibility in denying specific accounts. The administrator can use this module to reference a list of users who are not allowed to log in. Below is an example of how the module is used for the `vsftpd` FTP server in the `/etc/pam.d/vsftpd` PAM configuration file (the `\` character at the end of the first line in the following example is *not* necessary if the directive is on one line):

```
auth required /lib/security/pam_listfile.so item=user \  
sense=deny file=/etc/vsftpd.ftpusers onerr=succeed
```

This instructs PAM to consult the `/etc/vsftpd.ftpusers` file and deny access to the service for any listed user. The administrator can change the name of this file, and can keep separate lists for each service or use one central list to deny access to multiple services.

If the administrator wants to deny access to multiple services, a similar line can be added to the PAM configuration files, such as `/etc/pam.d/pop` and `/etc/pam.d/imap` for mail clients, or `/etc/pam.d/ssh` for SSH clients.

For more information about PAM, refer to [Section 2.5, “Pluggable Authentication Modules \(PAM\)”](#).

2.2.4.3. Limiting Root Access

Rather than completely denying access to the root user, the administrator may want to allow access only via setuid programs, such as `su` or `sudo`.

2.2.4.3.1. The su Command

When a user executes the `su` command, they are prompted for the root password and, after authentication, is given a root shell prompt.

Once logged in via the `su` command, the user *is* the root user and has absolute administrative access to the system³. In addition, once a user has become root, it is possible for them to use the `su` command to change to any other user on the system without being prompted for a password.

Because this program is so powerful, administrators within an organization may wish to limit who has access to the command.

One of the simplest ways to do this is to add users to the special administrative group called `wheel`. To do this, type the following command as root:

```
usermod -G wheel <username>
```

In the previous command, replace `<username>` with the username you want to add to the `wheel` group.

³This access is still subject to the restrictions imposed by SELinux, if it is enabled.

You can also use the **User Manager** to modify group memberships, as follows. Note: you need Administrator privileges to perform this procedure.

1. Click the **System** menu on the Panel, point to **Administration** and then click **Users and Groups** to display the User Manager. Alternatively, type the command `system-config-users` at a shell prompt.
2. Click the **Users** tab, and select the required user in the list of users.
3. Click **Properties** on the toolbar to display the User Properties dialog box (or choose **Properties** on the **File** menu).
4. Click the **Groups** tab, select the check box for the wheel group, and then click **OK**. Refer to [Figure 2.2, "Adding users to the "wheel" group."](#)
5. Open the PAM configuration file for `su (/etc/pam.d/su)` in a text editor and remove the comment `#` from the following line:

```
auth required /lib/security/$ISA/pam_wheel.so use_uid
```

This change means that only members of the administrative group `wheel` can use this program.

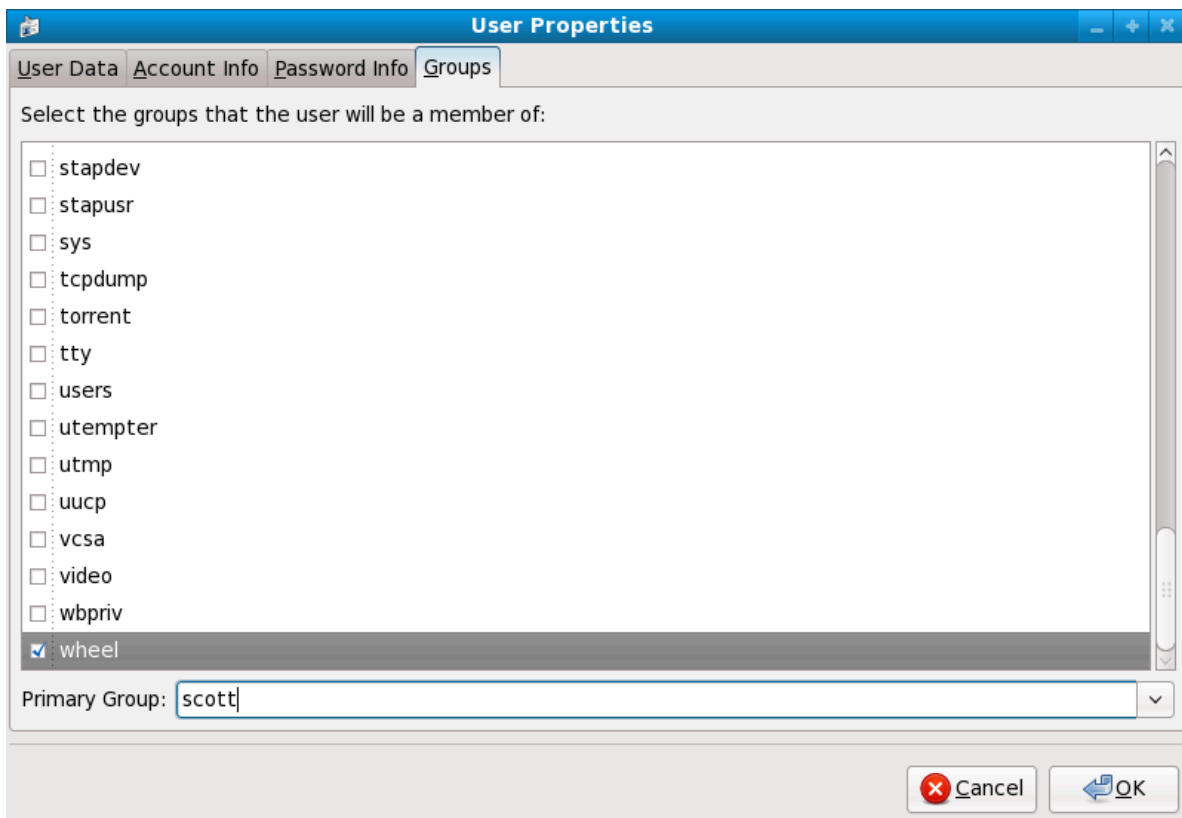


Figure 2.2. Adding users to the "wheel" group.

Note

The root user is part of the `wheel` group by default.

2.2.4.3.2. The `sudo` Command

The **sudo** command offers another approach to giving users administrative access. When trusted users precede an administrative command with **sudo**, they are prompted for *their own* password. Then, when they have been authenticated and assuming that the command is permitted, the administrative command is executed as if they were the root user.

The basic format of the **sudo** command is as follows:

```
sudo <command>
```

In the above example, *<command>* would be replaced by a command normally reserved for the root user, such as **mount**.



Important

Users of the **sudo** command should take extra care to log out before walking away from their machines since sudoers can use the command again without being asked for a password within a five minute period. This setting can be altered via the configuration file, `/etc/sudoers`.

The **sudo** command allows for a high degree of flexibility. For instance, only users listed in the `/etc/sudoers` configuration file are allowed to use the **sudo** command and the command is executed in *the user's* shell, not a root shell. This means the root shell can be completely disabled, as shown in [Section 2.2.4.2.1, "Disabling the Root Shell"](#).

The **sudo** command also provides a comprehensive audit trail. Each successful authentication is logged to the file `/var/log/messages` and the command issued along with the issuer's user name is logged to the file `/var/log/secure`.

Another advantage of the **sudo** command is that an administrator can allow different users access to specific commands based on their needs.

Administrators wanting to edit the **sudo** configuration file, `/etc/sudoers`, should use the **visudo** command.

To give someone full administrative privileges, type **visudo** and add a line similar to the following in the user privilege specification section:

```
juan ALL=(ALL) ALL
```

This example states that the user, `juan`, can use **sudo** from any host and execute any command.

The example below illustrates the granularity possible when configuring **sudo**:

```
%users localhost=/sbin/shutdown -h now
```

This example states that any user can issue the command `/sbin/shutdown -h now` as long as it is issued from the console.

The man page for **sudoers** has a detailed listing of options for this file.

2.2.5. Available Network Services

While user access to administrative controls is an important issue for system administrators within an organization, monitoring which network services are active is of paramount importance to anyone who administers and operates a Linux system.

Many services under Fedora behave as network servers. If a network service is running on a machine, then a server application (called a *daemon*), is listening for connections on one or more network ports. Each of these servers should be treated as a potential avenue of attack.

2.2.5.1. Risks To Services

Network services can pose many risks for Linux systems. Below is a list of some of the primary issues:

- *Denial of Service Attacks (DoS)* — By flooding a service with requests, a denial of service attack can render a system unusable as it tries to log and answer each request.
- *Distributed Denial of Service Attack (DDoS)* — A type of DoS attack which uses multiple compromised machines (often numbering in the thousands or more) to direct a co-ordinated attack on a service, flooding it with requests and making it unusable.
- *Script Vulnerability Attacks* — If a server is using scripts to execute server-side actions, as Web servers commonly do, a cracker can attack improperly written scripts. These script vulnerability attacks can lead to a buffer overflow condition or allow the attacker to alter files on the system.
- *Buffer Overflow Attacks* — Services that connect to ports numbered 0 through 1023 must run as an administrative user. If the application has an exploitable buffer overflow, an attacker could gain access to the system as the user running the daemon. Because exploitable buffer overflows exist, crackers use automated tools to identify systems with vulnerabilities, and once they have gained access, they use automated rootkits to maintain their access to the system.



Note

The threat of buffer overflow vulnerabilities is mitigated in Fedora by *ExecShield*, an executable memory segmentation and protection technology supported by x86-compatible uni- and multi-processor kernels. ExecShield reduces the risk of buffer overflow by separating virtual memory into executable and non-executable segments. Any program code that tries to execute outside of the executable segment (such as malicious code injected from a buffer overflow exploit) triggers a segmentation fault and terminates.

Execshield also includes support for *No eXecute (NX)* technology on AMD64 platforms and *eXecute Disable (XD)* technology on Itanium and Intel® 64 systems. These technologies work in conjunction with ExecShield to prevent malicious code from running in the executable portion of virtual memory with a granularity of 4KB of executable code, lowering the risk of attack from stealthy buffer overflow exploits.



Important

To limit exposure to attacks over the network, all services that are unused should be turned off.

2.2.5.2. Identifying and Configuring Services

To enhance security, most network services installed with Fedora are turned off by default. There are, however, some notable exceptions:

- **cupsd** — The default print server for Fedora.
- **lpd** — An alternative print server.
- **xinetd** — A super server that controls connections to a range of subordinate servers, such as **gssftp** and **telnet**.
- **sendmail** — The Sendmail *Mail Transport Agent* (MTA) is enabled by default, but only listens for connections from the localhost.
- **sshd** — The OpenSSH server, which is a secure replacement for Telnet.

When determining whether to leave these services running, it is best to use common sense and err on the side of caution. For example, if a printer is not available, do not leave **cupsd** running. The same is true for **portmap**. If you do not mount NFSv3 volumes or use NIS (the **ypbind** service), then **portmap** should be disabled.

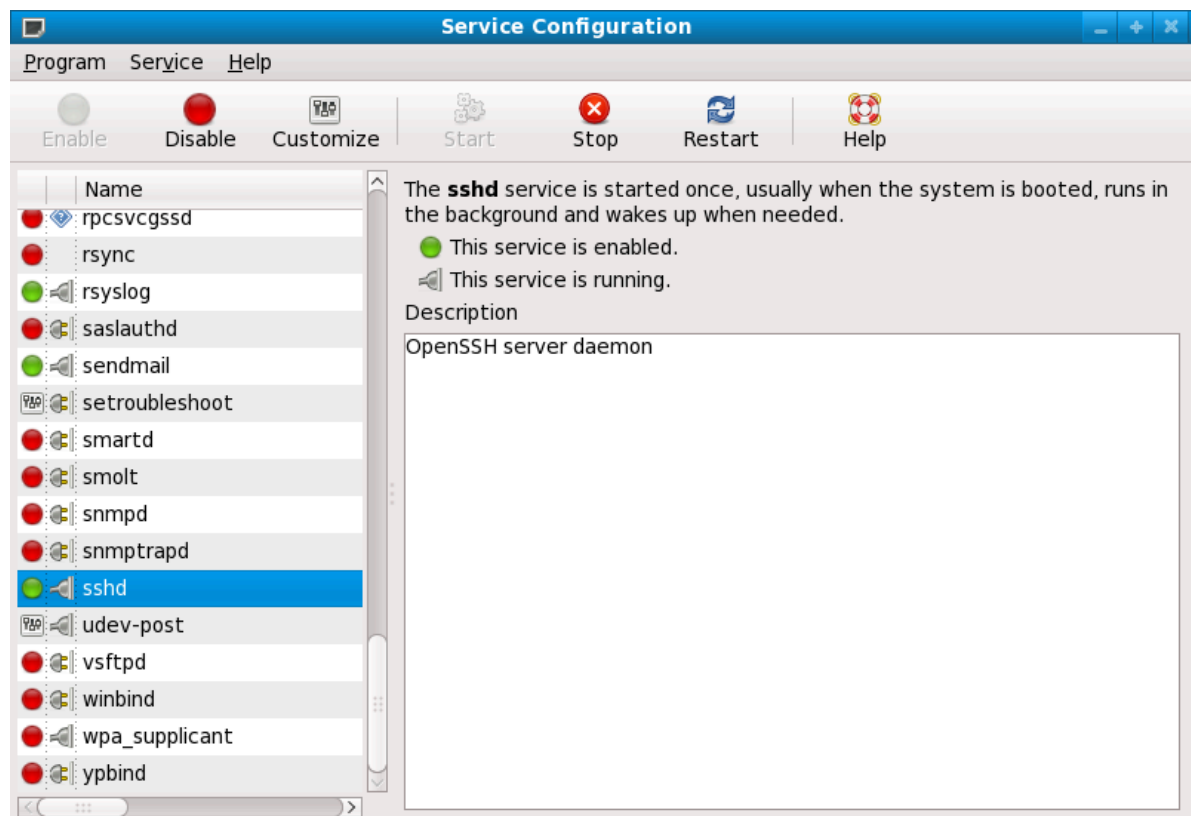


Figure 2.3. Services Configuration Tool

If unsure of the purpose for a particular service, the **Services Configuration Tool** has a description field, illustrated in [Figure 2.3, “Services Configuration Tool”](#), that provides additional information.

Checking which network services are available to start at boot time is only part of the story. You should also check which ports are open and listening. Refer to [Section 2.3.8, “Verifying Which Ports Are Listening”](#) for more information.

2.2.5.3. Insecure Services

Potentially, any network service is insecure. This is why turning off unused services is so important. Exploits for services are routinely revealed and patched, making it very important to regularly update packages associated with any network service. Refer to [Section 1.5, “Security Updates”](#) for more information.

Some network protocols are inherently more insecure than others. These include any services that:

- *Transmit Usernames and Passwords Over a Network Unencrypted* — Many older protocols, such as Telnet and FTP, do not encrypt the authentication session and should be avoided whenever possible.
- *Transmit Sensitive Data Over a Network Unencrypted* — Many protocols transmit data over the network unencrypted. These protocols include Telnet, FTP, HTTP, and SMTP. Many network file systems, such as NFS and SMB, also transmit information over the network unencrypted. It is the user's responsibility when using these protocols to limit what type of data is transmitted.

Remote memory dump services, like **netdump**, transmit the contents of memory over the network unencrypted. Memory dumps can contain passwords or, even worse, database entries and other sensitive information.

Other services like **finger** and **rwhod** reveal information about users of the system.

Examples of inherently insecure services include **rlogin**, **rsh**, **telnet**, and **vsftpd**.

All remote login and shell programs (**rlogin**, **rsh**, and **telnet**) should be avoided in favor of SSH. Refer to [Section 2.2.7, “Security Enhanced Communication Tools”](#) for more information about **sshd**.

FTP is not as inherently dangerous to the security of the system as remote shells, but FTP servers must be carefully configured and monitored to avoid problems. Refer to [Section 2.3.6, “Securing FTP”](#) for more information about securing FTP servers.

Services that should be carefully implemented and behind a firewall include:

- **finger**
- **authd** (this was called **identd** in previous Fedora releases.)
- **netdump**
- **netdump-server**
- **nfs**
- **rwhod**
- **sendmail**
- **smb** (Samba)
- **yppasswdd**
- **ypserv**
- **ypxfrd**

More information on securing network services is available in [Section 2.3, “Server Security”](#).

The next section discusses tools available to set up a simple firewall.

2.2.6. Personal Firewalls

After the *necessary* network services are configured, it is important to implement a firewall.



Important

You should configure the necessary services and implement a firewall *before* connecting to the Internet or any other network that you do not trust.

Firewalls prevent network packets from accessing the system's network interface. If a request is made to a port that is blocked by a firewall, the request is ignored. If a service is listening on one of these blocked ports, it does not receive the packets and is effectively disabled. For this reason, care should be taken when configuring a firewall to block access to ports not in use, while not blocking access to ports used by configured services.

For most users, the best tool for configuring a simple firewall is the graphical firewall configuration tool which ships with Fedora: the **Firewall Configuration Tool** (`system-config-firewall`). This tool creates broad **iptables** rules for a general-purpose firewall using a control panel interface.

Refer to [Section 2.9.2, “Basic Firewall Configuration”](#) for more information about using this application and its available options.

For advanced users and server administrators, manually configuring a firewall with **iptables** is probably a better option. Refer to [Section 2.9, “Firewalls”](#) for more information. Refer to [Section 2.10, “IPTables”](#) for a comprehensive guide to the **iptables** command.

2.2.7. Security Enhanced Communication Tools

As the size and popularity of the Internet has grown, so has the threat of communication interception. Over the years, tools have been developed to encrypt communications as they are transferred over the network.

Fedora ships with two basic tools that use high-level, public-key-cryptography-based encryption algorithms to protect information as it travels over the network.

- *OpenSSH* — A free implementation of the SSH protocol for encrypting network communication.
- *Gnu Privacy Guard (GPG)* — A free implementation of the PGP (Pretty Good Privacy) encryption application for encrypting data.

OpenSSH is a safer way to access a remote machine and replaces older, unencrypted services like **telnet** and **rsh**. OpenSSH includes a network service called **sshd** and three command line client applications:

- **ssh** — A secure remote console access client.
- **scp** — A secure remote copy command.
- **sftp** — A secure pseudo-ftp client that allows interactive file transfer sessions.

Refer to [Section 3.6, “Secure Shell”](#) for more information regarding OpenSSH.



Important

Although the `sshd` service is inherently secure, the service *must* be kept up-to-date to prevent security threats. Refer to *Section 1.5, “Security Updates”* for more information.

GPG is one way to ensure private email communication. It can be used both to email sensitive data over public networks and to protect sensitive data on hard drives.

2.3. Server Security

When a system is used as a server on a public network, it becomes a target for attacks. Hardening the system and locking down services is therefore of paramount importance for the system administrator.

Before delving into specific issues, review the following general tips for enhancing server security:

- Keep all services current, to protect against the latest threats.
- Use secure protocols whenever possible.
- Serve only one type of network service per machine whenever possible.
- Monitor all servers carefully for suspicious activity.

2.3.1. Securing Services With TCP Wrappers and `xinetd`

TCP Wrappers provide access control to a variety of services. Most modern network services, such as SSH, Telnet, and FTP, make use of TCP Wrappers, which stand guard between an incoming request and the requested service.

The benefits offered by TCP Wrappers are enhanced when used in conjunction with `xinetd`, a super server that provides additional access, logging, binding, redirection, and resource utilization control.



Note

It is a good idea to use iptables firewall rules in conjunction with TCP Wrappers and `xinetd` to create redundancy within service access controls. Refer to *Section 2.9, “Firewalls”* for more information about implementing firewalls with iptables commands.

The following subsections assume a basic knowledge of each topic and focus on specific security options.

2.3.1.1. Enhancing Security With TCP Wrappers

TCP Wrappers are capable of much more than denying access to services. This section illustrates how they can be used to send connection banners, warn of attacks from particular hosts, and enhance logging functionality. Refer to the `hosts_options` man page for information about the TCP Wrapper functionality and control language.

2.3.1.1.1. TCP Wrappers and Connection Banners

Displaying a suitable banner when users connect to a service is a good way to let potential attackers know that the system administrator is being vigilant. You can also control what information about the

system is presented to users. To implement a TCP Wrappers banner for a service, use the `banner` option.

This example implements a banner for `vsftpd`. To begin, create a banner file. It can be anywhere on the system, but it must have same name as the daemon. For this example, the file is called `/etc/banners/vsftpd` and contains the following line:

```
220-Hello, %c
220-All activity on ftp.example.com is logged.
220-Inappropriate use will result in your access privileges being removed.
```

The `%c` token supplies a variety of client information, such as the username and hostname, or the username and IP address to make the connection even more intimidating.

For this banner to be displayed to incoming connections, add the following line to the `/etc/hosts.allow` file:

```
vsftpd : ALL : banners /etc/banners/
```

2.3.1.1.2. TCP Wrappers and Attack Warnings

If a particular host or network has been detected attacking the server, TCP Wrappers can be used to warn the administrator of subsequent attacks from that host or network using the `spawn` directive.

In this example, assume that a cracker from the 206.182.68.0/24 network has been detected attempting to attack the server. Place the following line in the `/etc/hosts.deny` file to deny any connection attempts from that network, and to log the attempts to a special file:

```
ALL : 206.182.68.0 : spawn /bin/ 'date' %c %d >> /var/log/intruder_alert
```

The `%d` token supplies the name of the service that the attacker was trying to access.

To allow the connection and log it, place the `spawn` directive in the `/etc/hosts.allow` file.



Note

Because the `spawn` directive executes any shell command, it is a good idea to create a special script to notify the administrator or execute a chain of commands in the event that a particular client attempts to connect to the server.

2.3.1.1.3. TCP Wrappers and Enhanced Logging

If certain types of connections are of more concern than others, the log level can be elevated for that service using the `severity` option.

For this example, assume that anyone attempting to connect to port 23 (the Telnet port) on an FTP server is a cracker. To denote this, place an `emerg` flag in the log files instead of the default flag, `info`, and deny the connection.

To do this, place the following line in `/etc/hosts.deny`:

```
in.telnetd : ALL : severity emerg
```

This uses the default **authpriv** logging facility, but elevates the priority from the default value of **info** to **emerg**, which posts log messages directly to the console.

2.3.1.2. Enhancing Security With xinetd

This section focuses on using **xinetd** to set a trap service and using it to control resource levels available to any given **xinetd** service. Setting resource limits for services can help thwart *Denial of Service* (DoS) attacks. Refer to the man pages for **xinetd** and **xinetd.conf** for a list of available options.

2.3.1.2.1. Setting a Trap

One important feature of **xinetd** is its ability to add hosts to a global **no_access** list. Hosts on this list are denied subsequent connections to services managed by **xinetd** for a specified period or until **xinetd** is restarted. You can do this using the **SENSOR** attribute. This is an easy way to block hosts attempting to scan the ports on the server.

The first step in setting up a **SENSOR** is to choose a service you do not plan on using. For this example, Telnet is used.

Edit the file `/etc/xinetd.d/telnet` and change the `flags` line to read:

```
flags          = SENSOR
```

Add the following line:

```
deny_time     = 30
```

This denies any further connection attempts to that port by that host for 30 minutes. Other acceptable values for the **deny_time** attribute are **FOREVER**, which keeps the ban in effect until **xinetd** is restarted, and **NEVER**, which allows the connection and logs it.

Finally, the last line should read:

```
disable       = no
```

This enables the trap itself.

While using **SENSOR** is a good way to detect and stop connections from undesirable hosts, it has two drawbacks:

- It does not work against stealth scans.
- An attacker who knows that a **SENSOR** is running can mount a Denial of Service attack against particular hosts by forging their IP addresses and connecting to the forbidden port.

2.3.1.2.2. Controlling Server Resources

Another important feature of **xinetd** is its ability to set resource limits for services under its control.

It does this using the following directives:


- `cps = <number_of_connections> <wait_period>` — Limits the rate of incoming connections. This directive takes two arguments:

- `<number_of_connections>` — The number of connections per second to handle. If the rate of incoming connections is higher than this, the service is temporarily disabled. The default value is fifty (50).
- `<wait_period>` — The number of seconds to wait before re-enabling the service after it has been disabled. The default interval is ten (10) seconds.
- `instances = <number_of_connections>` — Specifies the total number of connections allowed to a service. This directive accepts either an integer value or **UNLIMITED**.
- `per_source = <number_of_connections>` — Specifies the number of connections allowed to a service by each host. This directive accepts either an integer value or **UNLIMITED**.
- `rlimit_as = <number[K|M]>` — Specifies the amount of memory address space the service can occupy in kilobytes or megabytes. This directive accepts either an integer value or **UNLIMITED**.
- `rlimit_cpu = <number_of_seconds>` — Specifies the amount of time in seconds that a service may occupy the CPU. This directive accepts either an integer value or **UNLIMITED**.

Using these directives can help prevent any single **xinetd** service from overwhelming the system, resulting in a denial of service.

2.3.2. Securing Portmap

The **portmap** service is a dynamic port assignment daemon for RPC services such as NIS and NFS. It has weak authentication mechanisms and has the ability to assign a wide range of ports for the services it controls. For these reasons, it is difficult to secure.



Note
Securing **portmap** only affects NFSv2 and NFSv3 implementations, since NFSv4 no longer requires it. If you plan to implement an NFSv2 or NFSv3 server, then **portmap** is required, and the following section applies.

If running RPC services, follow these basic rules.

2.3.2.1. Protect portmap With TCP Wrappers

It is important to use TCP Wrappers to limit which networks or hosts have access to the **portmap** service since it has no built-in form of authentication.

Further, use *only* IP addresses when limiting access to the service. Avoid using hostnames, as they can be forged by DNS poisoning and other methods.

2.3.2.2. Protect portmap With iptables

To further restrict access to the **portmap** service, it is a good idea to add iptables rules to the server and restrict access to specific networks.

Below are two example iptables commands. The first allows TCP connections to the port 111 (used by the **portmap** service) from the 192.168.0.0/24 network. The second allows TCP connections to the same port from the localhost. This is necessary for the **sgi_fam** service used by **Nautilus**. All other packets are dropped.

```
iptables -A INPUT -p tcp -s! 192.168.0.0/24 --dport 111 -j DROP
iptables -A INPUT -p tcp -s 127.0.0.1 --dport 111 -j ACCEPT
```

To similarly limit UDP traffic, use the following command.

```
iptables -A INPUT -p udp -s! 192.168.0.0/24 --dport 111 -j DROP
```



Note

Refer to *Section 2.9, “Firewalls”* for more information about implementing firewalls with iptables commands.

2.3.3. Securing NIS

The *Network Information Service* (NIS) is an RPC service, called **ypserv**, which is used in conjunction with **portmap** and other related services to distribute maps of usernames, passwords, and other sensitive information to any computer claiming to be within its domain.

An NIS server is comprised of several applications. They include the following:

- **/usr/sbin/rpc.yppasswdd** — Also called the **yppasswdd** service, this daemon allows users to change their NIS passwords.
- **/usr/sbin/rpc.ypxfrd** — Also called the **ypxfrd** service, this daemon is responsible for NIS map transfers over the network.
- **/usr/sbin/yppush** — This application propagates changed NIS databases to multiple NIS servers.
- **/usr/sbin/ypserv** — This is the NIS server daemon.

NIS is somewhat insecure by today's standards. It has no host authentication mechanisms and transmits all of its information over the network unencrypted, including password hashes. As a result, extreme care must be taken when setting up a network that uses NIS. This is further complicated by the fact that the default configuration of NIS is inherently insecure.

It is recommended that anyone planning to implement an NIS server first secure the **portmap** service as outlined in *Section 2.3.2, “Securing Portmap”*, then address the following issues, such as network planning.

2.3.3.1. Carefully Plan the Network

Because NIS transmits sensitive information unencrypted over the network, it is important the service be run behind a firewall and on a segmented and secure network. Whenever NIS information is transmitted over an insecure network, it risks being intercepted. Careful network design can help prevent severe security breaches.

2.3.3.2. Use a Password-like NIS Domain Name and Hostname

Any machine within an NIS domain can use commands to extract information from the server without authentication, as long as the user knows the NIS server's DNS hostname and NIS domain name.

For instance, if someone either connects a laptop computer into the network or breaks into the network from outside (and manages to spoof an internal IP address), the following command reveals the `/etc/passwd` map:

```
ypcat -d <NIS_domain> -h <DNS_hostname> passwd
```

If this attacker is a root user, they can obtain the `/etc/shadow` file by typing the following command:

```
ypcat -d <NIS_domain> -h <DNS_hostname> shadow
```



Note

If Kerberos is used, the `/etc/shadow` file is not stored within an NIS map.

To make access to NIS maps harder for an attacker, create a random string for the DNS hostname, such as `o7hfawtgmhwg.domain.com`. Similarly, create a *different* randomized NIS domain name. This makes it much more difficult for an attacker to access the NIS server.

2.3.3.3. Edit the `/var/yp/securenets` File

If the `/var/yp/securenets` file is blank or does not exist (as is the case after a default installation), NIS listens to all networks. One of the first things to do is to put netmask/network pairs in the file so that `ypserv` only responds to requests from the appropriate network.

Below is a sample entry from a `/var/yp/securenets` file:

```
255.255.255.0    192.168.0.0
```



Warning

Never start an NIS server for the first time without creating the `/var/yp/securenets` file.

This technique does not provide protection from an IP spoofing attack, but it does at least place limits on what networks the NIS server services.

2.3.3.4. Assign Static Ports and Use iptables Rules

All of the servers related to NIS can be assigned specific ports except for `rpc.yppasswdd` — the daemon that allows users to change their login passwords. Assigning ports to the other two NIS server daemons, `rpc.ypxfrd` and `ypserv`, allows for the creation of firewall rules to further protect the NIS server daemons from intruders.

To do this, add the following lines to `/etc/sysconfig/network`:

```
YPSERV_ARGS="-p 834" YPXFRD_ARGS="-p 835"
```

The following iptables rules can then be used to enforce which network the server listens to for these ports:

```
iptables -A INPUT -p ALL -s! 192.168.0.0/24 --dport 834 -j DROP
iptables -A INPUT -p ALL -s! 192.168.0.0/24 --dport 835 -j DROP
```

This means that the server only allows connections to ports 834 and 835 if the requests come from the 192.168.0.0/24 network, regardless of the protocol.



Note

Refer to [Section 2.9, “Firewalls”](#) for more information about implementing firewalls with iptables commands.

2.3.3.5. Use Kerberos Authentication

One of the issues to consider when NIS is used for authentication is that whenever a user logs into a machine, a password hash from the `/etc/shadow` map is sent over the network. If an intruder gains access to an NIS domain and sniffs network traffic, they can collect usernames and password hashes. With enough time, a password cracking program can guess weak passwords, and an attacker can gain access to a valid account on the network.

Since Kerberos uses secret-key cryptography, no password hashes are ever sent over the network, making the system far more secure. Refer to [Section 2.7, “Kerberos”](#) for more information about Kerberos.

2.3.4. Securing NFS



Important

The version of NFS included in Fedora, NFSv4, no longer requires the `portmap` service as outlined in [Section 2.3.2, “Securing Portmap”](#). NFS traffic now utilizes TCP in all versions, rather than UDP, and requires it when using NFSv4. NFSv4 now includes Kerberos user and group authentication, as part of the `RPCSEC_GSS` kernel module. Information on `portmap` is still included, since Fedora supports NFSv2 and NFSv3, both of which utilize `portmap`.

2.3.4.1. Carefully Plan the Network

Now that NFSv4 has the ability to pass all information encrypted using Kerberos over a network, it is important that the service be configured correctly if it is behind a firewall or on a segmented network. NFSv2 and NFSv3 still pass data insecurely, and this should be taken into consideration. Careful network design in all of these regards can help prevent security breaches.

2.3.4.2. Beware of Syntax Errors

The NFS server determines which file systems to export and which hosts to export these directories to by consulting the `/etc/exports` file. Be careful not to add extraneous spaces when editing this file.

For instance, the following line in the `/etc/exports` file shares the directory `/tmp/nfs/` to the host `bob.example.com` with read/write permissions.

```
/tmp/nfs/    bob.example.com(rw)
```


The following line in the `/etc/exports` file, on the other hand, shares the same directory to the host `bob.example.com` with read-only permissions and shares it to the *world* with read/write permissions due to a single space character after the hostname.

```
/tmp/nfs/    bob.example.com (rw)
```

It is good practice to check any configured NFS shares by using the `showmount` command to verify what is being shared:

```
showmount -e <hostname>
```

2.3.4.3. Do Not Use the `no_root_squash` Option

By default, NFS shares change the root user to the `nfsnobody` user, an unprivileged user account. This changes the owner of all root-created files to `nfsnobody`, which prevents uploading of programs with the `setuid` bit set.

If `no_root_squash` is used, remote root users are able to change any file on the shared file system and leave applications infected by trojans for other users to inadvertently execute.

2.3.4.4. NFS Firewall Configuration

The ports used for NFS are assigned dynamically by `rpcbind`, which can cause problems when creating firewall rules. To simplify this process, use the `/etc/sysconfig/nfs` file to specify which ports are to be used:

- **MOUNTD_PORT** — TCP and UDP port for `mountd` (`rpc.mountd`)
- **STATD_PORT** — TCP and UDP port for `status` (`rpc.statd`)
- **LOCKD_TCP** — TCP port for `nlockmgr` (`rpc.lockd`)
- **LOCKD_UDP** — UDP port `nlockmgr` (`rpc.lockd`)

Port numbers specified must not be used by any other service. Configure your firewall to allow the port numbers specified, as well as TCP and UDP port 2049 (NFS).

Run the `rpcinfo -p` command on the NFS server to see which ports and RPC programs are being used.

2.3.5. Securing the Apache HTTP Server

The Apache HTTP Server is one of the most stable and secure services that ships with Fedora. A large number of options and techniques are available to secure the Apache HTTP Server — too numerous to delve into deeply here. The following section briefly explains good practices when running the Apache HTTP Server.

Always verify that any scripts running on the system work as intended *before* putting them into production. Also, ensure that only the root user has write permissions to any directory containing scripts or CGIs. To do this, run the following commands as the root user:

1.

```
chown root <directory_name>
```

2. `chmod 755 <directory_name>`

System administrators should be careful when using the following configuration options (configured in `/etc/httpd/conf/httpd.conf`):

FollowSymLinks

This directive is enabled by default, so be sure to use caution when creating symbolic links to the document root of the Web server. For instance, it is a bad idea to provide a symbolic link to `/`.

Indexes

This directive is enabled by default, but may not be desirable. To prevent visitors from browsing files on the server, remove this directive.

UserDir

The `UserDir` directive is disabled by default because it can confirm the presence of a user account on the system. To enable user directory browsing on the server, use the following directives:

```
UserDir enabled
UserDir disabled root
```

These directives activate user directory browsing for all user directories other than `/root/`. To add users to the list of disabled accounts, add a space-delimited list of users on the `UserDir disabled` line.



Important

Do not remove the `IncludesNoExec` directive. By default, the *Server-Side Includes* (SSI) module cannot execute commands. It is recommended that you do not change this setting unless absolutely necessary, as it could, potentially, enable an attacker to execute commands on the system.

2.3.6. Securing FTP

The *File Transfer Protocol* (FTP) is an older TCP protocol designed to transfer files over a network. Because all transactions with the server, including user authentication, are unencrypted, it is considered an insecure protocol and should be carefully configured.

Fedora provides three FTP servers.

- **gssftpd** — A Kerberos-aware `xinetd`-based FTP daemon that does not transmit authentication information over the network.
- **Red Hat Content Accelerator (tux)** — A kernel-space Web server with FTP capabilities.
- **vsftpd** — A standalone, security oriented implementation of the FTP service.

The following security guidelines are for setting up the **vsftpd** FTP service.

2.3.6.1. FTP Greeting Banner

Before submitting a username and password, all users are presented with a greeting banner. By default, this banner includes version information useful to crackers trying to identify weaknesses in a system.

To change the greeting banner for **vsftpd**, add the following directive to the **/etc/vsftpd/vsftpd.conf** file:

```
ftpd_banner=<insert_greeting_here>
```

Replace *<insert_greeting_here>* in the above directive with the text of the greeting message.

For mutli-line banners, it is best to use a banner file. To simplify management of multiple banners, place all banners in a new directory called **/etc/banners/**. The banner file for FTP connections in this example is **/etc/banners/ftp.msg**. Below is an example of what such a file may look like:

```
##### # Hello, all activity on ftp.example.com is logged. #####
```



Note

It is not necessary to begin each line of the file with **220** as specified in [Section 2.3.1.1.1, “TCP Wrappers and Connection Banners”](#).

To reference this greeting banner file for **vsftpd**, add the following directive to the **/etc/vsftpd/vsftpd.conf** file:

```
banner_file=/etc/banners/ftp.msg
```

It also is possible to send additional banners to incoming connections using TCP Wrappers as described in [Section 2.3.1.1.1, “TCP Wrappers and Connection Banners”](#).

2.3.6.2. Anonymous Access

The presence of the **/var/ftp/** directory activates the anonymous account.

The easiest way to create this directory is to install the **vsftpd** package. This package establishes a directory tree for anonymous users and configures the permissions on directories to read-only for anonymous users.

By default the anonymous user cannot write to any directories.



Warning

If enabling anonymous access to an FTP server, be aware of where sensitive data is stored.

2.3.6.2.1. Anonymous Upload

To allow anonymous users to upload files, it is recommended that a write-only directory be created within **/var/ftp/pub/**.

To do this, type the following command:

```
mkdir /var/ftp/pub/upload
```

Next, change the permissions so that anonymous users cannot view the contents of the directory:

```
chmod 730 /var/ftp/pub/upload
```

A long format listing of the directory should look like this:

```
drwx-wx---  2 root    ftp          4096 Feb 13 20:05 upload
```



Warning

Administrators who allow anonymous users to read and write in directories often find that their servers become a repository of stolen software.

Additionally, under **vsftpd**, add the following line to the `/etc/vsftpd/vsftpd.conf` file:

```
anon_upload_enable=YES
```

2.3.6.3. User Accounts

Because FTP transmits unencrypted usernames and passwords over insecure networks for authentication, it is a good idea to deny system users access to the server from their user accounts.

To disable all user accounts in **vsftpd**, add the following directive to `/etc/vsftpd/vsftpd.conf`:

```
local_enable=NO
```

2.3.6.3.1. Restricting User Accounts

To disable FTP access for specific accounts or specific groups of accounts, such as the root user and those with **sudo** privileges, the easiest way is to use a PAM list file as described in [Section 2.2.4.2.4, “Disabling Root Using PAM”](#). The PAM configuration file for **vsftpd** is `/etc/pam.d/vsftpd`.

It is also possible to disable user accounts within each service directly.

To disable specific user accounts in **vsftpd**, add the username to `/etc/vsftpd.ftpusers`

2.3.6.4. Use TCP Wrappers To Control Access

Use TCP Wrappers to control access to either FTP daemon as outlined in [Section 2.3.1.1, “Enhancing Security With TCP Wrappers”](#).

2.3.7. Securing Sendmail

Sendmail is a Mail Transfer Agent (MTA) that uses the Simple Mail Transfer Protocol (SMTP) to deliver electronic messages between other MTAs and to email clients or delivery agents. Although many MTAs are capable of encrypting traffic between one another, most do not, so sending email over any public networks is considered an inherently insecure form of communication.

It is recommended that anyone planning to implement a Sendmail server address the following issues.

2.3.7.1. Limiting a Denial of Service Attack

Because of the nature of email, a determined attacker can flood the server with mail fairly easily and cause a denial of service. By setting limits to the following directives in `/etc/mail/sendmail.mc`, the effectiveness of such attacks is limited.

- **confCONNECTION_RATE_THROTTLE** — The number of connections the server can receive per second. By default, Sendmail does not limit the number of connections. If a limit is set and reached, further connections are delayed.
- **confMAX_DAEMON_CHILDREN** — The maximum number of child processes that can be spawned by the server. By default, Sendmail does not assign a limit to the number of child processes. If a limit is set and reached, further connections are delayed.
- **confMIN_FREE_BLOCKS** — The minimum number of free blocks which must be available for the server to accept mail. The default is 100 blocks.
- **confMAX_HEADERS_LENGTH** — The maximum acceptable size (in bytes) for a message header.
- **confMAX_MESSAGE_SIZE** — The maximum acceptable size (in bytes) for a single message.

2.3.7.2. NFS and Sendmail

Never put the mail spool directory, `/var/spool/mail/`, on an NFS shared volume.

Because NFSv2 and NFSv3 do not maintain control over user and group IDs, two or more users can have the same UID, and receive and read each other's mail.



Note

With NFSv4 using Kerberos, this is not the case, since the **SECRPC_GSS** kernel module does not utilize UID-based authentication. However, it is still considered good practice *not* to put the mail spool directory on NFS shared volumes.

2.3.7.3. Mail-only Users

To help prevent local user exploits on the Sendmail server, it is best for mail users to only access the Sendmail server using an email program. Shell accounts on the mail server should not be allowed and all user shells in the `/etc/passwd` file should be set to `/sbin/nologin` (with the possible exception of the root user).

2.3.8. Verifying Which Ports Are Listening

After configuring network services, it is important to pay attention to which ports are actually listening on the system's network interfaces. Any open ports can be evidence of an intrusion.

There are two basic approaches for listing the ports that are listening on the network. The less reliable approach is to query the network stack using commands such as **netstat -an** or **lsof -i**. This method is less reliable since these programs do not connect to the machine from the network, but rather check to see what is running on the system. For this reason, these applications are frequent

targets for replacement by attackers. Crackers attempt to cover their tracks if they open unauthorized network ports by replacing **netstat** and **lsof** with their own, modified versions.

A more reliable way to check which ports are listening on the network is to use a port scanner such as **nmap**.

The following command issued from the console determines which ports are listening for TCP connections from the network:

```
nmap -sT -O localhost
```

The output of this command appears as follows:

```
Starting Nmap 4.68 ( http://nmap.org ) at 2009-03-06 12:08 EST
Interesting ports on localhost.localdomain (127.0.0.1):
Not shown: 1711 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
25/tcp    open  smtp
111/tcp   open  rpcbind
113/tcp   open  auth
631/tcp   open  ipp
834/tcp   open  unknown
2601/tcp  open  zebra
32774/tcp open  sometimes-rpc11
Device type: general purpose
Running: Linux 2.6.X
OS details: Linux 2.6.17 - 2.6.24
Uptime: 4.122 days (since Mon Mar  2 09:12:31 2009)
Network Distance: 0 hops
OS detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 1.420 seconds
```

This output shows the system is running **portmap** due to the presence of the `sunrpc` service. However, there is also a mystery service on port 834. To check if the port is associated with the official list of known services, type:

```
cat /etc/services | grep 834
```

This command returns no output. This indicates that while the port is in the reserved range (meaning 0 through 1023) and requires root access to open, it is not associated with a known service.

Next, check for information about the port using **netstat** or **lsof**. To check for port 834 using **netstat**, use the following command:

```
netstat -anp | grep 834
```

The command returns the following output:

```
tcp    0      0 0.0.0.0:834      0.0.0.0:*        LISTEN  653/ybind
```

The presence of the open port in **netstat** is reassuring because a cracker opening a port surreptitiously on a hacked system is not likely to allow it to be revealed through this command. Also, the `[p]` option reveals the process ID (PID) of the service that opened the port. In this case, the open port belongs to **ybind** (NIS), which is an RPC service handled in conjunction with the **portmap** service.

The **lsof** command reveals similar information to **netstat** since it is also capable of linking open ports to services:

```
lsof -i | grep 834
```

The relevant portion of the output from this command follows:

ybind	653	0	7u	IPv4	1319	TCP *:834 (LISTEN)
ybind	655	0	7u	IPv4	1319	TCP *:834 (LISTEN)
ybind	656	0	7u	IPv4	1319	TCP *:834 (LISTEN)
ybind	657	0	7u	IPv4	1319	TCP *:834 (LISTEN)

These tools reveal a great deal about the status of the services running on a machine. These tools are flexible and can provide a wealth of information about network services and configuration. Refer to the man pages for **lsof**, **netstat**, **nmap**, and **services** for more information.

2.4. Single Sign-on (SSO)

2.4.1. Introduction

The Fedora SSO functionality reduces the number of times Fedora desktop users have to enter their passwords. Several major applications leverage the same underlying authentication and authorization mechanisms so that users can log in to Fedora from the log-in screen, and then not need to re-enter their passwords. These applications are detailed below.

In addition, users can log in to their machines even when there is no network (*offline mode*) or where network connectivity is unreliable, for example, wireless access. In the latter case, services will degrade gracefully.

2.4.1.1. Supported Applications

The following applications are currently supported by the unified log-in scheme in Fedora:

- Login
- Screensaver
- Firefox and Thunderbird

2.4.1.2. Supported Authentication Mechanisms

Fedora currently supports the following authentication mechanisms:

- Kerberos name/password login
- Smart card/PIN login

2.4.1.3. Supported Smart Cards

Fedora has been tested with the Cyberflex e-gate card and reader, but any card that complies with both Java card 2.1.1 and Global Platform 2.0.1 specifications should operate correctly, as should any reader that is supported by PCSC-lite.

Fedora has also been tested with Common Access Cards (CAC). The supported reader for CAC is the SCM SCR 331 USB Reader.

As of Fedora 5.2, Gemalto smart cards (Cyberflex Access 64k v2, standard with DER SHA1 value configured as in PKCSI v2.1) are now supported. These smart cards now use readers compliant with Chip/Smart Card Interface Devices (CCID).

2.4.1.4. Advantages of Fedora Single Sign-on

Numerous security mechanisms currently exist that utilize a large number of protocols and credential stores. Examples include SSL, SSH, IPsec, and Kerberos. Fedora SSO aims to unify these schemes to support the requirements listed above. This does not mean replacing Kerberos with X.509v3 certificates, but rather uniting them to reduce the burden on both system users and the administrators who manage them.

To achieve this goal, Fedora:

- Provides a single, shared instance of the NSS crypto libraries on each operating system.
- Ships the Certificate System's Enterprise Security Client (ESC) with the base operating system. The ESC application monitors smart card insertion events. If it detects that the user has inserted a smart card that was designed to be used with the Fedora Certificate System server product, it displays a user interface instructing the user how to enroll that smart card.
- Unifies Kerberos and NSS so that users who log in to the operating system using a smart card also obtain a Kerberos credential (which allows them to log in to file servers, etc.)

2.4.2. Getting Started with your new Smart Card

Before you can use your smart card to log in to your system and take advantage of the increased security options this technology provides, you need to perform some basic installation and configuration steps. These are described below.



Note

This section provides a high-level view of getting started with your smart card. More detailed information is available in the Red Hat Certificate System Enterprise Security Client Guide.

1. Log in with your Kerberos name and password
2. Make sure you have the **nss-tools** package loaded.
3. Download and install your corporate-specific root certificates. Use the following command to install the root CA certificate:

```
certutil -A -d /etc/pki/nssdb -n "root ca cert" -t "CT,C,C" -i ./ca_cert_in_base64_format.crt
```

4. Verify that you have the following RPMs installed on your system: `esc`, `pam_pkcs11`, `coolkey`, `ifd-egate`, `ccid`, `gdm`, `authconfig`, and `authconfig-gtk`.
5. Enable Smart Card Login Support

- a. On the Gnome Title Bar, select System->Administration->Authentication.
 - b. Type your machine's root password if necessary.
 - c. In the Authentication Configuration dialog, click the **Authentication** tab.
 - d. Select the **Enable Smart Card Support** check box.
 - e. Click the **Configure Smart Card...** button to display the Smartcard Settings dialog, and specify the required settings:
 - **Require smart card for login** — Clear this check box. After you have successfully logged in with the smart card you can select this option to prevent users from logging in without a smart card.
 - **Card Removal Action** — This controls what happens when you remove the smart card after you have logged in. The available options are:
 - **Lock** — Removing the smart card locks the X screen.
 - **Ignore** — Removing the smart card has no effect.
6. If you need to enable the Online Certificate Status Protocol (OCSP), open the `/etc/pam_pkcs11/pam_pkcs11.conf` file, and locate the following line:

enable_ocsp = false;

Change this value to true, as follows:

enable_ocsp = true;
 7. Enroll your smart card
 8. If you are using a CAC card, you also need to perform the following steps:
 - a. Change to the root account and create a file called `/etc/pam_pkcs11/cn_map`.
 - b. Add the following entry to the `cn_map` file:

`MY.CAC_CN.123454 -> myloginid`

where `MY.CAC_CN.123454` is the Common Name on your CAC and `myloginid` is your UNIX login ID.
 9. Logout

2.4.2.1. Troubleshooting

If you have trouble getting your smart card to work, try using the following command to locate the source of the problem:

```
pklogin_finder debug
```

If you run the `pklogin_finder` tool in debug mode while an enrolled smart card is plugged in, it attempts to output information about the validity of certificates, and if it is successful in attempting to map a login ID from the certificates that are on the card.

2.4.3. How Smart Card Enrollment Works

Smart cards are said to be *enrolled* when they have received an appropriate certificate signed by a valid Certificate Authority (CA). This involves several steps, described below:

1. The user inserts their smart card into the smart card reader on their workstation. This event is recognized by the Enterprise Security Client (ESC).
2. The enrollment page is displayed on the user's desktop. The user completes the required details and the user's system then connects to the Token Processing System (TPS) and the CA.
3. The TPS enrolls the smart card using a certificate signed by the CA.

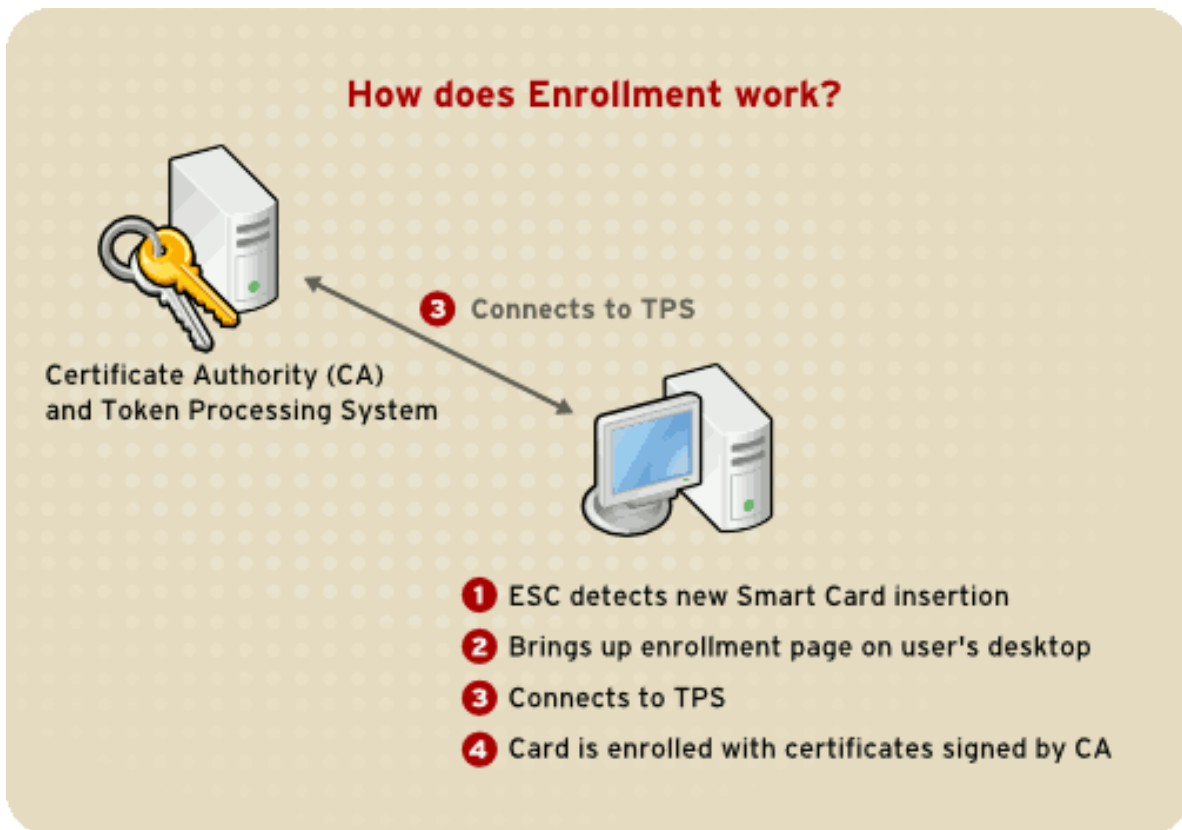


Figure 2.4. How Smart Card Enrollment Works

2.4.4. How Smart Card Login Works

This section provides a brief overview of the process of logging in using a smart card.

1. When the user inserts their smart card into the smart card reader, this event is recognized by the PAM facility, which prompts for the user's PIN.
2. The system then looks up the user's current certificates and verifies their validity. The certificate is then mapped to the user's UID.

3. This is validated against the KDC and login granted.

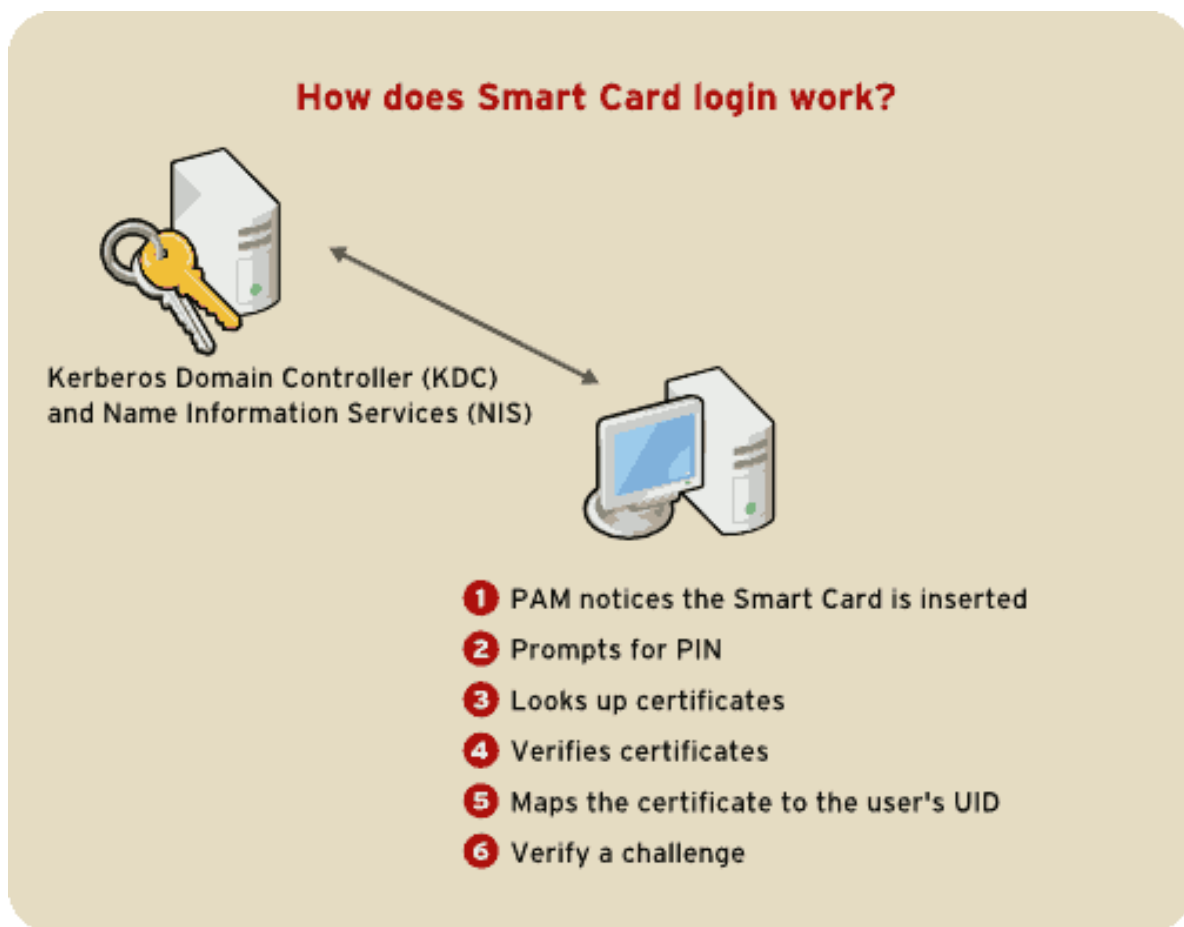


Figure 2.5. How Smart Card Login Works



Note

You cannot log in with a card that has not been enrolled, even if it has been formatted. You need to log in with a formatted, enrolled card, or not using a smart card, before you can enroll a new card.

Refer to [Section 2.7, “Kerberos”](#) and [Section 2.5, “Pluggable Authentication Modules \(PAM\)”](#) for more information on Kerberos and PAM.

2.4.5. Configuring Firefox to use Kerberos for SSO

You can configure Firefox to use Kerberos for Single Sign-on. In order for this functionality to work correctly, you need to configure your web browser to send your Kerberos credentials to the appropriate KDC. The following section describes the configuration changes and other requirements to achieve this.

1. In the address bar of Firefox, type **about:config** to display the list of current configuration options.
2. In the **Filter** field, type **negotiate** to restrict the list of options.

3. Double-click the `network.negotiate-auth.trusted-uris` entry to display the *Enter string value* dialog box.
4. Enter the name of the domain against which you want to authenticate, for example, `.example.com`.
5. Repeat the above procedure for the `network.negotiate-auth.delegation-uris` entry, using the same domain.

Note

You can leave this value blank, as it allows Kerberos ticket passing, which is not required.

If you do not see these two configuration options listed, your version of Firefox may be too old to support Negotiate authentication, and you should consider upgrading.

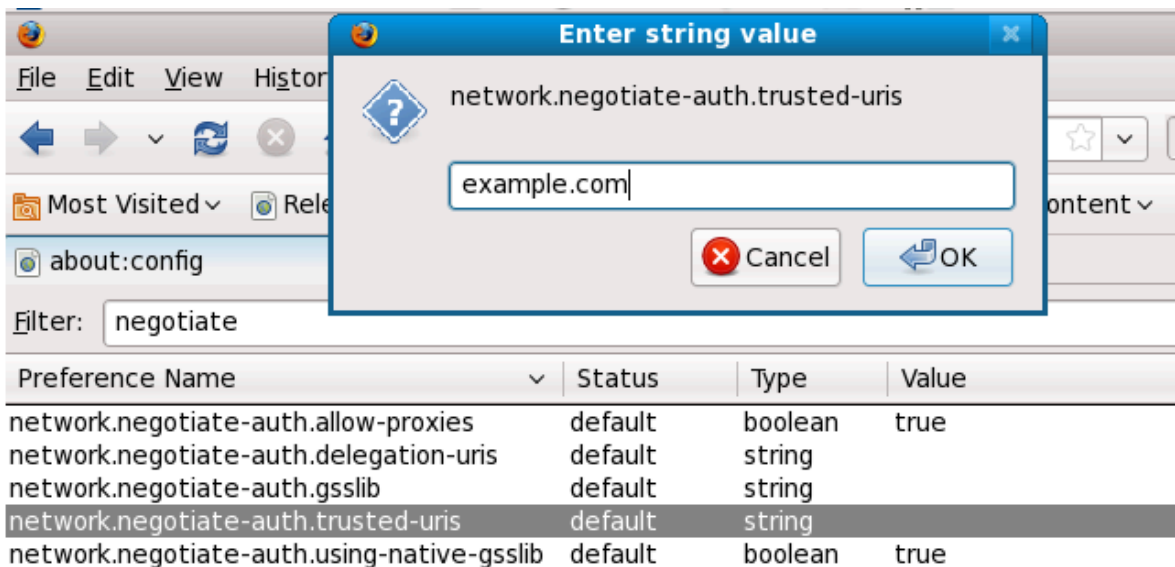


Figure 2.6. Configuring Firefox for SSO with Kerberos

You now need to ensure that you have Kerberos tickets. In a command shell, type `kinit` to retrieve Kerberos tickets. To display the list of available tickets, type `klist`. The following shows an example output from these commands:

```
[user@host ~] $ kinit
Password for user@EXAMPLE.COM:

[user@host ~] $ klist
Ticket cache: FILE:/tmp/krb5cc_10920
Default principal: user@EXAMPLE.COM

Valid starting    Expires          Service principal
10/26/06 23:47:54 10/27/06 09:47:54  krbtgt/USER.COM@USER.COM
        renew until 10/26/06 23:47:54

Kerberos 4 ticket cache: /tmp/tkt10920
klist: You have no tickets cached
```

2.4.5.1. Troubleshooting

If you have followed the configuration steps above and Negotiate authentication is not working, you can turn on verbose logging of the authentication process. This could help you find the cause of the problem. To enable verbose logging, use the following procedure:

1. Close all instances of Firefox.
2. Open a command shell, and enter the following commands:

```
export NSPR_LOG_MODULES=negotiateauth:5
export NSPR_LOG_FILE=/tmp/moz.log
```

3. Restart Firefox *from that shell*, and visit the website you were unable to authenticate to earlier. Information will be logged to `/tmp/moz.log`, and may give a clue to the problem. For example:

```
-1208550944[90039d0]: entering nsNegotiateAuth::GetNextToken()
-1208550944[90039d0]: gss_init_sec_context() failed: Miscellaneous failure
No credentials cache found
```

This indicates that you do not have Kerberos tickets, and need to run **kinit**.

If you are able to run **kinit** successfully from your machine but you are unable to authenticate, you might see something like this in the log file:

```
-1208994096[8d683d8]: entering nsAuthGSSAPI::GetNextToken()
-1208994096[8d683d8]: gss_init_sec_context() failed: Miscellaneous failure
Server not found in Kerberos database
```

This generally indicates a Kerberos configuration problem. Make sure that you have the correct entries in the `[domain_realm]` section of the `/etc/krb5.conf` file. For example:

```
.example.com = EXAMPLE.COM
example.com = EXAMPLE.COM
```

If nothing appears in the log it is possible that you are behind a proxy, and that proxy is stripping off the HTTP headers required for Negotiate authentication. As a workaround, you can try to connect to the server using HTTPS instead, which allows the request to pass through unmodified. Then proceed to debug using the log file, as described above.

2.5. Pluggable Authentication Modules (PAM)

Programs that grant users access to a system use *authentication* to verify each other's identity (that is, to establish that a user is who they say they are).

Historically, each program had its own way of authenticating users. In Fedora, many programs are configured to use a centralized authentication mechanism called *Pluggable Authentication Modules* (PAM).

PAM uses a pluggable, modular architecture, which affords the system administrator a great deal of flexibility in setting authentication policies for the system.

In most situations, the default PAM configuration file for a PAM-aware application is sufficient. Sometimes, however, it is necessary to edit a PAM configuration file. Because misconfiguration of PAM can compromise system security, it is important to understand the structure of these files

before making any modifications. Refer to [Section 2.5.3, “PAM Configuration File Format”](#) for more information.

2.5.1. Advantages of PAM

PAM offers the following advantages:

- a common authentication scheme that can be used with a wide variety of applications.
- significant flexibility and control over authentication for both system administrators and application developers.
- a single, fully-documented library which allows developers to write programs without having to create their own authentication schemes.

2.5.2. PAM Configuration Files

The `/etc/pam.d/` directory contains the PAM configuration files for each PAM-aware application. In earlier versions of PAM, the `/etc/pam.conf` file was used, but this file is now deprecated and is only used if the `/etc/pam.d/` directory does not exist.

2.5.2.1. PAM Service Files

Each PAM-aware application or *service* has a file in the `/etc/pam.d/` directory. Each file in this directory has the same name as the service to which it controls access.

The PAM-aware program is responsible for defining its service name and installing its own PAM configuration file in the `/etc/pam.d/` directory. For example, the `login` program defines its service name as `login` and installs the `/etc/pam.d/login` PAM configuration file.

2.5.3. PAM Configuration File Format

Each PAM configuration file contains a group of directives formatted as follows:

```
<module interface> <control flag> <module name> <module arguments>
```

Each of these elements is explained in the following sections.

2.5.3.1. Module Interface

Four types of PAM module interface are currently available. Each of these corresponds to a different aspect of the authorization process:

- **auth** — This module interface authenticates use. For example, it requests and verifies the validity of a password. Modules with this interface can also set credentials, such as group memberships or Kerberos tickets.
- **account** — This module interface verifies that access is allowed. For example, it may check if a user account has expired or if a user is allowed to log in at a particular time of day.
- **password** — This module interface is used for changing user passwords.
- **session** — This module interface configures and manages user sessions. Modules with this interface can also perform additional tasks that are needed to allow access, like mounting a user's home directory and making the user's mailbox available.



Note

An individual module can provide any or all module interfaces. For instance, `pam_unix.so` provides all four module interfaces.

In a PAM configuration file, the module interface is the first field defined. For example, a typical line in a configuration may look like this:

```
auth required pam_unix.so
```

This instructs PAM to use the `pam_unix.so` module's `auth` interface.

2.5.3.1.1. Stacking Module Interfaces

Module interface directives can be *stacked*, or placed upon one another, so that multiple modules are used together for one purpose. If a module's control flag uses the "sufficient" or "requisite" value (refer to [Section 2.5.3.2, "Control Flag"](#) for more information on these flags), then the order in which the modules are listed is important to the authentication process.

Stacking makes it easy for an administrator to require specific conditions to exist before allowing the user to authenticate. For example, the `reboot` command normally uses several stacked modules, as seen in its PAM configuration file:

```
[root@MyServer ~]# cat /etc/pam.d/reboot
#%PAM-1.0
auth sufficient pam_rootok.so
auth required pam_console.so
#auth include system-auth
account required pam_permit.so
```

- The first line is a comment and is not processed.
- **auth sufficient pam_rootok.so** — This line uses the `pam_rootok.so` module to check whether the current user is root, by verifying that their UID is 0. If this test succeeds, no other modules are consulted and the command is executed. If this test fails, the next module is consulted.
- **auth required pam_console.so** — This line uses the `pam_console.so` module to attempt to authenticate the user. If this user is already logged in at the console, `pam_console.so` checks whether there is a file in the `/etc/security/console.apps/` directory with the same name as the service name (`reboot`). If such a file exists, authentication succeeds and control is passed to the next module.
- **#auth include system-auth** — This line is commented and is not processed.
- **account required pam_permit.so** — This line uses the `pam_permit.so` module to allow the root user or anyone logged in at the console to reboot the system.

2.5.3.2. Control Flag

All PAM modules generate a success or failure result when called. Control flags tell PAM what to do with the result. Modules can be stacked in a particular order, and the control flags determine how important the success or failure of a particular module is to the overall goal of authenticating the user to the service.

There are four predefined control flags:

- **required** — The module result must be successful for authentication to continue. If the test fails at this point, the user is not notified until the results of all module tests that reference that interface are complete.
- **requisite** — The module result must be successful for authentication to continue. However, if a test fails at this point, the user is notified immediately with a message reflecting the first failed **required** or **requisite** module test.
- **sufficient** — The module result is ignored if it fails. However, if the result of a module flagged **sufficient** is successful *and* no previous modules flagged **required** have failed, then no other results are required and the user is authenticated to the service.
- **optional** — The module result is ignored. A module flagged as **optional** only becomes necessary for successful authentication when no other modules reference the interface.



Important

The order in which **required** modules are called is not critical. Only the **sufficient** and **requisite** control flags cause order to become important.

A newer control flag syntax that allows for more precise control is now available for PAM.

The **pam.d** man page, and the PAM documentation, located in the `/usr/share/doc/pam-<version-number>/` directory, where `<version-number>` is the version number for PAM on your system, describe this newer syntax in detail.

2.5.3.3. Module Name

The module name provides PAM with the name of the pluggable module containing the specified module interface. In older versions of Fedora, the full path to the module was provided in the PAM configuration file. However, since the advent of *multilib* systems, which store 64-bit PAM modules in the `/lib64/security/` directory, the directory name is omitted because the application is linked to the appropriate version of **libpam**, which can locate the correct version of the module.

2.5.3.4. Module Arguments

PAM uses *arguments* to pass information to a pluggable module during authentication for some modules.

For example, the **pam_userdb.so** module uses information stored in a Berkeley DB file to authenticate the user. Berkeley DB is an open source database system embedded in many applications. The module takes a **db** argument so that Berkeley DB knows which database to use for the requested service.

The following is a typical **pam_userdb.so** line in a PAM configuration. The `<path-to-file>` is the full path to the Berkeley DB database file:

```
auth required pam_userdb.so db=<path-to-file>
```


Invalid arguments are *generally* ignored and do not otherwise affect the success or failure of the PAM module. Some modules, however, may fail on invalid arguments. Most modules report errors to the `/var/log/secure` file.

2.5.4. Sample PAM Configuration Files

The following is a sample PAM application configuration file:

```
##%PAM-1.0
auth required pam_securetty.so
auth required pam_unix.so nullok
auth required pam_nologin.so
account required pam_unix.so
password required pam_cracklib.so retry=3
password required pam_unix.so shadow nullok use_authok
session required pam_unix.so
```

- The first line is a comment, indicated by the hash mark (#) at the beginning of the line.
- Lines two through four stack three modules for login authentication.

auth required pam_securetty.so — This module ensures that *if* the user is trying to log in as root, the tty on which the user is logging in is listed in the `/etc/securetty` file, *if* that file exists.

If the tty is not listed in the file, any attempt to log in as root fails with a `Login incorrect` message.

auth required pam_unix.so nullok — This module prompts the user for a password and then checks the password using the information stored in `/etc/passwd` and, if it exists, `/etc/shadow`.

- The argument **nullok** instructs the **pam_unix.so** module to allow a blank password.
- **auth required pam_nologin.so** — This is the final authentication step. It checks whether the `/etc/nologin` file exists. If it exists and the user is not root, authentication fails.



Note

In this example, all three **auth** modules are checked, even if the first **auth** module fails. This prevents the user from knowing at what stage their authentication failed. Such knowledge in the hands of an attacker could allow them to more easily deduce how to crack the system.

- **account required pam_unix.so** — This module performs any necessary account verification. For example, if shadow passwords have been enabled, the account interface of the **pam_unix.so** module checks to see if the account has expired or if the user has not changed the password within the allowed grace period.
- **password required pam_cracklib.so retry=3** — If a password has expired, the password component of the **pam_cracklib.so** module prompts for a new password. It then tests the newly created password to see whether it can easily be determined by a dictionary-based password cracking program.

- The argument **retry=3** specifies that if the test fails the first time, the user has two more chances to create a strong password.
- **password required pam_unix.so shadow nullok use_authtok** — This line specifies that if the program changes the user's password, it should use the **password** interface of the **pam_unix.so** module to do so.
 - The argument **shadow** instructs the module to create shadow passwords when updating a user's password.
 - The argument **nullok** instructs the module to allow the user to change their password *from* a blank password, otherwise a null password is treated as an account lock.
 - The final argument on this line, **use_authtok**, provides a good example of the importance of order when stacking PAM modules. This argument instructs the module not to prompt the user for a new password. Instead, it accepts any password that was recorded by a previous password module. In this way, all new passwords must pass the **pam_cracklib.so** test for secure passwords before being accepted.
- **session required pam_unix.so** — The final line instructs the session interface of the **pam_unix.so** module to manage the session. This module logs the user name and the service type to **/var/log/secure** at the beginning and end of each session. This module can be supplemented by stacking it with other session modules for additional functionality.

2.5.5. Creating PAM Modules

You can create or add new PAM modules at any time for use by PAM-aware applications.

For example, a developer might create a one-time-password creation method and write a PAM module to support it. PAM-aware programs can immediately use the new module and password method without being recompiled or otherwise modified.

This allows developers and system administrators to mix-and-match, as well as test, authentication methods for different programs without recompiling them.

Documentation on writing modules is included in the **/usr/share/doc/pam-<version-number>/** directory, where **<version-number>** is the version number for PAM on your system.

2.5.6. PAM and Administrative Credential Caching

A number of graphical administrative tools in Fedora provide users with elevated privileges for up to five minutes using the **pam_timestamp.so** module. It is important to understand how this mechanism works, because a user who walks away from a terminal while **pam_timestamp.so** is in effect leaves the machine open to manipulation by anyone with physical access to the console.

In the PAM timestamp scheme, the graphical administrative application prompts the user for the root password when it is launched. When the user has been authenticated, the **pam_timestamp.so** module creates a timestamp file. By default, this is created in the **/var/run/sudo/** directory. If the timestamp file already exists, graphical administrative programs do not prompt for a password. Instead, the **pam_timestamp.so** module freshens the timestamp file, reserving an extra five minutes of unchallenged administrative access for the user.

You can verify the actual state of the timestamp file by inspecting the `/var/run/sudo/<user>` file. For the desktop, the relevant file is **unknown:root**. If it is present and its timestamp is less than five minutes old, the credentials are valid.

The existence of the timestamp file is indicated by an authentication icon, which appears in the notification area of the panel.



Figure 2.7. The Authentication Icon

2.5.6.1. Removing the Timestamp File

Before abandoning a console where a PAM timestamp is active, it is recommended that the timestamp file be destroyed. To do this from a graphical environment, click the authentication icon on the panel. This causes a dialog box to appear. Click the **Forget Authorization** button to destroy the active timestamp file.

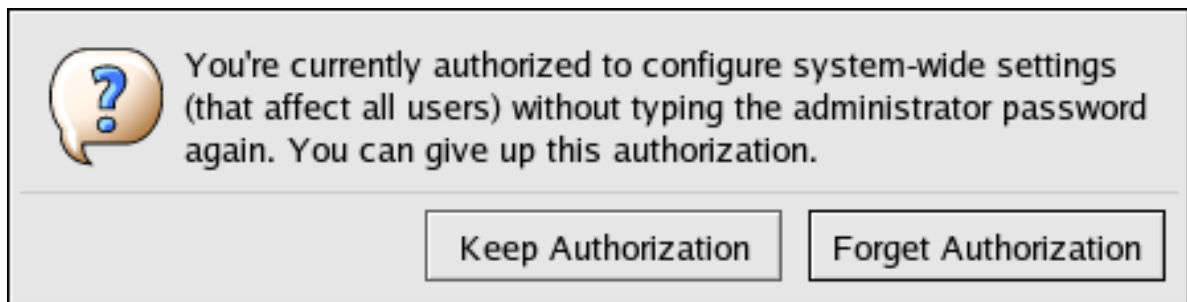


Figure 2.8. Dismiss Authentication Dialog

You should be aware of the following with respect to the PAM timestamp file:

- If logged in to the system remotely using **ssh**, use the `/sbin/pam_timestamp_check -k root` command to destroy the timestamp file.
- You need to run the `/sbin/pam_timestamp_check -k root` command from the same terminal window from which you launched the privileged application.
- You must be logged in as the user who originally invoked the `pam_timestamp.so` module in order to use the `/sbin/pam_timestamp_check -k` command. Do not log in as root to use this command.
- If you want to kill the credentials on the desktop (without using the **Forget Authorization** action on the icon), use the following command:

```
/sbin/pam_timestamp_check -k root </dev/null >/dev/null 2>/dev/null
```

Failure to use this command will only remove the credentials (if any) from the pty where you run the command.

Refer to the `pam_timestamp_check` man page for more information about destroying the timestamp file using `pam_timestamp_check`.

2.5.6.2. Common pam_timestamp Directives

The `pam_timestamp.so` module accepts several directives. The following are the two most commonly used options:

- `timestamp_timeout` — Specifies the period (in seconds) for which the timestamp file is valid. The default value is 300 (five minutes).
- `timestampdir` — Specifies the directory in which the timestamp file is stored. The default value is `/var/run/sudo/`.

Refer to [Section 2.9.9.1, “Installed Firewall Documentation”](#) for more information about controlling the `pam_timestamp.so` module.

2.5.7. PAM and Device Ownership

In Fedora, the first user who logs in at the physical console of the machine can manipulate certain devices and perform certain tasks normally reserved for the root user. This is controlled by a PAM module called `pam_console.so`.

2.5.7.1. Device Ownership

When a user logs in to a Fedora system, the `pam_console.so` module is called by `login` or the graphical login programs, `gdm`, `kdm`, and `xdm`. If this user is the first user to log in at the physical console — referred to as the *console user* — the module grants the user ownership of a variety of devices normally owned by root. The console user owns these devices until the last local session for that user ends. After this user has logged out, ownership of the devices reverts back to the root user.

The devices affected include, but are not limited to, sound cards, diskette drives, and CD-ROM drives.

This facility allows a local user to manipulate these devices without obtaining root access, thus simplifying common tasks for the console user.

You can modify the list of devices controlled by `pam_console.so` by editing the following files:

- `/etc/security/console.perms`
- `/etc/security/console.perms.d/50-default.perms`

You can change the permissions of different devices than those listed in the above files, or override the specified defaults. Rather than modify the `50-default.perms` file, you should create a new file (for example, `xx-name.perms`) and enter the required modifications. The name of the new default file must begin with a number higher than 50 (for example, `51-default.perms`). This will override the defaults in the `50-default.perms` file.



Warning

If the `gdm`, `kdm`, or `xdm` display manager configuration file has been altered to allow remote users to log in *and* the host is configured to run at runlevel 5, it is advisable to change the `<console>` and `<xconsole>` directives in the `/etc/security/console.perms` to the following values:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]* :0\.[0-9] :0
<xconsole>=:0\.[0-9] :0
```

This prevents remote users from gaining access to devices and restricted applications on the machine.

If the **gdm**, **kdm**, or **xdm** display manager configuration file has been altered to allow remote users to log in *and* the host is configured to run at any multiple user runlevel other than 5, it is advisable to remove the **<xconsole>** directive entirely and change the **<console>** directive to the following value:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]*
```

2.5.7.2. Application Access

The console user also has access to certain programs configured for use in the **/etc/security/console.apps/** directory.

This directory contains configuration files which enable the console user to run certain applications in **/sbin** and **/usr/sbin**.

These configuration files have the same name as the applications that they set up.

One notable group of applications that the console user has access to are three programs that shut down or reboot the system:

- **/sbin/halt**
- **/sbin/reboot**
- **/sbin/poweroff**

Because these are PAM-aware applications, they call the **pam_console.so** module as a requirement for use.

Refer to [Section 2.9.9.1, “Installed Firewall Documentation”](#) for more information.

2.5.8. Additional Resources

The following resources further explain methods to use and configure PAM. In addition to these resources, read the PAM configuration files on the system to better understand how they are structured.

2.5.8.1. Installed PAM Documentation

- PAM-related man pages — Several man pages exist for the various applications and configuration files involved with PAM. The following is a list of some of the more important man pages.

Configuration Files

- **pam** — Good introductory information on PAM, including the structure and purpose of the PAM configuration files.

Note that this man page discusses both **/etc/pam.conf** and individual configuration files in the **/etc/pam.d/** directory. By default, Fedora uses the individual configuration files in the **/etc/pam.d/** directory, ignoring **/etc/pam.conf** even if it exists.

- **pam_console** — Describes the purpose of the **pam_console.so** module. It also describes the appropriate syntax for an entry within a PAM configuration file.
- **console.apps** — Describes the format and options available in the **/etc/security/console.apps** configuration file, which defines which applications are accessible by the console user assigned by PAM.
- **console.perms** — Describes the format and options available in the **/etc/security/console.perms** configuration file, which specifies the console user permissions assigned by PAM.
- **pam_timestamp** — Describes the **pam_timestamp.so** module.
- **/usr/share/doc/pam-<version-number>** — Contains a *System Administrators' Guide*, a *Module Writers' Manual*, and the *Application Developers' Manual*, as well as a copy of the PAM standard, DCE-RFC 86.0, where **<version-number>** is the version number of PAM.
- **/usr/share/doc/pam-<version-number>/txts/README.pam_timestamp** — Contains information about the **pam_timestamp.so** PAM module, where **<version-number>** is the version number of PAM.

2.5.8.2. Useful PAM Websites

- <http://www.kernel.org/pub/linux/libs/pam/> — The primary distribution website for the Linux-PAM project, containing information on various PAM modules, a FAQ, and additional PAM documentation.



Note

The documentation in the above website is for the last released upstream version of PAM and might not be 100% accurate for the PAM version included in Fedora.

2.6. TCP Wrappers and xinetd

Controlling access to network services is one of the most important security tasks facing a server administrator. Fedora provides several tools for this purpose. For example, an **iptables**-based firewall filters out unwelcome network packets within the kernel's network stack. For network services that utilize it, *TCP Wrappers* add an additional layer of protection by defining which hosts are or are not allowed to connect to "wrapped" network services. One such wrapped network service is the *xinetd super server*. This service is called a super server because it controls connections to a subset of network services and further refines access control.

Figure 2.9, "Access Control to Network Services" is a basic illustration of how these tools work together to protect network services.

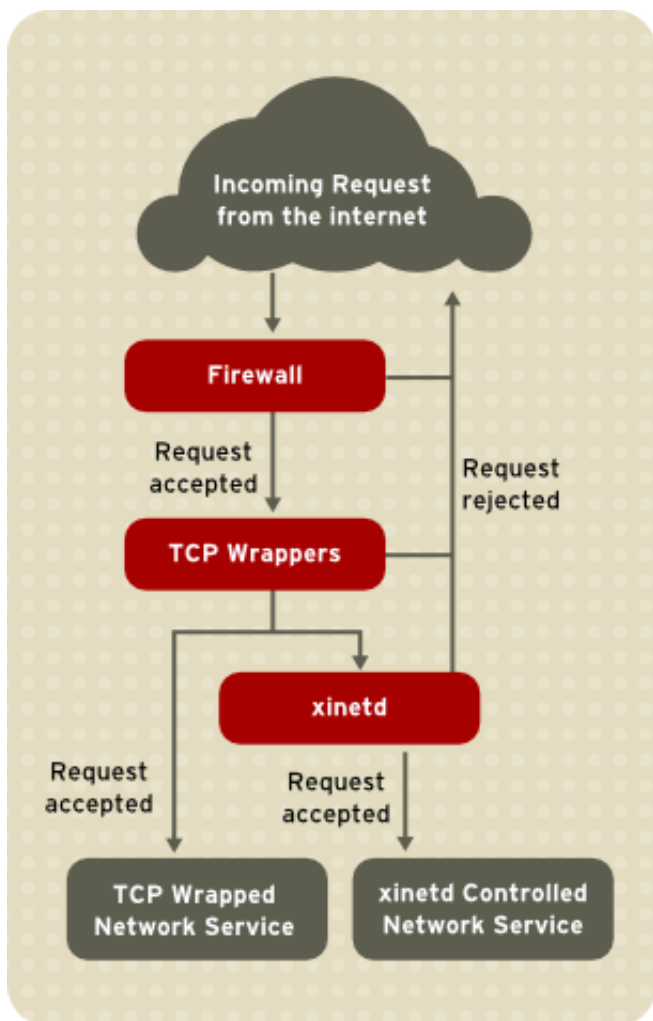


Figure 2.9. Access Control to Network Services

This chapter focuses on the role of TCP Wrappers and `xinetd` in controlling access to network services and reviews how these tools can be used to enhance both logging and utilization management. Refer to [Section 2.10, “IPTables”](#) for information about using firewalls with `iptables`.

2.6.1. TCP Wrappers

The TCP Wrappers package (`tcp_wrappers`) is installed by default and provides host-based access control to network services. The most important component within the package is the `/usr/lib/libwrap.a` library. In general terms, a TCP-wrapped service is one that has been compiled against the `libwrap.a` library.

When a connection attempt is made to a TCP-wrapped service, the service first references the host's access files (`/etc/hosts.allow` and `/etc/hosts.deny`) to determine whether or not the client is allowed to connect. In most cases, it then uses the syslog daemon (`syslogd`) to write the name of the requesting client and the requested service to `/var/log/secure` or `/var/log/messages`.

If a client is allowed to connect, TCP Wrappers release control of the connection to the requested service and take no further part in the communication between the client and the server.

In addition to access control and logging, TCP Wrappers can execute commands to interact with the client before denying or releasing control of the connection to the requested network service.

Because TCP Wrappers are a valuable addition to any server administrator's arsenal of security tools, most network services within Fedora are linked to the **libwrap.a** library. Some such applications include `/usr/sbin/sshd`, `/usr/sbin/sendmail`, and `/usr/sbin/xinetd`.

Note

To determine if a network service binary is linked to **libwrap.a**, type the following command as the root user:

```
ldd <binary-name> | grep libwrap
```

Replace *<binary-name>* with the name of the network service binary.

If the command returns straight to the prompt with no output, then the network service is *not* linked to **libwrap.a**.

The following example indicates that `/usr/sbin/sshd` is linked to **libwrap.a**:

```
[root@myServer ~]# ldd /usr/sbin/sshd | grep libwrap
    libwrap.so.0 => /lib/libwrap.so.0 (0x00655000)
[root@myServer ~]#
```

2.6.1.1. Advantages of TCP Wrappers

TCP Wrappers provide the following advantages over other network service control techniques:

- *Transparency to both the client and the wrapped network service* — Both the connecting client and the wrapped network service are unaware that TCP Wrappers are in use. Legitimate users are logged and connected to the requested service while connections from banned clients fail.
- *Centralized management of multiple protocols* — TCP Wrappers operate separately from the network services they protect, allowing many server applications to share a common set of access control configuration files, making for simpler management.

2.6.2. TCP Wrappers Configuration Files

To determine if a client is allowed to connect to a service, TCP Wrappers reference the following two files, which are commonly referred to as *hosts access* files:

- `/etc/hosts.allow`
- `/etc/hosts.deny`

When a TCP-wrapped service receives a client request, it performs the following steps:

1. *It references `/etc/hosts.allow`.* — The TCP-wrapped service sequentially parses the `/etc/hosts.allow` file and applies the first rule specified for that service. If it finds a matching rule, it allows the connection. If not, it moves on to the next step.
2. *It references `/etc/hosts.deny`.* — The TCP-wrapped service sequentially parses the `/etc/hosts.deny` file. If it finds a matching rule, it denies the connection. If not, it grants access to the service.

The following are important points to consider when using TCP Wrappers to protect network services:

- Because access rules in **hosts.allow** are applied first, they take precedence over rules specified in **hosts.deny**. Therefore, if access to a service is allowed in **hosts.allow**, a rule denying access to that same service in **hosts.deny** is ignored.
- The rules in each file are read from the top down and the first matching rule for a given service is the only one applied. The order of the rules is extremely important.
- If no rules for the service are found in either file, or if neither file exists, access to the service is granted.
- TCP-wrapped services do not cache the rules from the hosts access files, so any changes to **hosts.allow** or **hosts.deny** take effect immediately, without restarting network services.



Warning

If the last line of a hosts access file is not a newline character (created by pressing the **Enter** key), the last rule in the file fails and an error is logged to either `/var/log/messages` or `/var/log/secure`. This is also the case for a rule that spans multiple lines without using the backslash character. The following example illustrates the relevant portion of a log message for a rule failure due to either of these circumstances:

```
warning: /etc/hosts.allow, line 20: missing newline or line too long
```

2.6.2.1. Formatting Access Rules

The format for both `/etc/hosts.allow` and `/etc/hosts.deny` is identical. Each rule must be on its own line. Blank lines or lines that start with a hash (#) are ignored.

Each rule uses the following basic format to control access to network services:

```
<daemon list>: <client list> [: <option>: <option>: ...]
```

- `<daemon list>` — A comma-separated list of process names (*not* service names) or the ALL wildcard. The daemon list also accepts operators (refer to [Section 2.6.2.1.4, “Operators”](#)) to allow greater flexibility.
- `<client list>` — A comma-separated list of hostnames, host IP addresses, special patterns, or wildcards which identify the hosts affected by the rule. The client list also accepts operators listed in [Section 2.6.2.1.4, “Operators”](#) to allow greater flexibility.
- `<option>` — An optional action or colon-separated list of actions performed when the rule is triggered. Option fields support expansions, launch shell commands, allow or deny access, and alter logging behavior.



Note

More information on the specialist terms above can be found elsewhere in this Guide:

- [Section 2.6.2.1.1, “Wildcards”](#)

- [Section 2.6.2.1.2, “Patterns”](#)
- [Section 2.6.2.2.4, “Expansions”](#)
- [Section 2.6.2.2, “Option Fields”](#)

The following is a basic sample hosts access rule:

```
vsftpd : .example.com
```

This rule instructs TCP Wrappers to watch for connections to the FTP daemon (`vsftpd`) from any host in the `example.com` domain. If this rule appears in **hosts.allow**, the connection is accepted. If this rule appears in **hosts.deny**, the connection is rejected.

The next sample hosts access rule is more complex and uses two option fields:

```
sshd : .example.com \ : spawn /bin/echo `/bin/date` access denied>>/var/log/sshd.log \ : deny
```

Note that each option field is preceded by the backslash (`\`). Use of the backslash prevents failure of the rule due to length.

This sample rule states that if a connection to the SSH daemon (`sshd`) is attempted from a host in the `example.com` domain, execute the **echo** command to append the attempt to a special log file, and deny the connection. Because the optional **deny** directive is used, this line denies access even if it appears in the **hosts.allow** file. Refer to [Section 2.6.2.2, “Option Fields”](#) for a more detailed look at available options.

2.6.2.1.1. Wildcards

Wildcards allow TCP Wrappers to more easily match groups of daemons or hosts. They are used most frequently in the client list field of access rules.

The following wildcards are available:

- ALL — Matches everything. It can be used for both the daemon list and the client list.
- LOCAL — Matches any host that does not contain a period (`.`), such as `localhost`.
- KNOWN — Matches any host where the hostname and host address are known or where the user is known.
- UNKNOWN — Matches any host where the hostname or host address are unknown or where the user is unknown.
- PARANOID — Matches any host where the hostname does not match the host address.



Important

The KNOWN, UNKNOWN, and PARANOID wildcards should be used with care, because they rely on functioning DNS server for correct operation. Any disruption to name resolution may prevent legitimate users from gaining access to a service.

2.6.2.1.2. Patterns

Patterns can be used in the client field of access rules to more precisely specify groups of client hosts.

The following is a list of common patterns for entries in the client field:

- *Hostname beginning with a period (.)* — Placing a period at the beginning of a hostname matches all hosts sharing the listed components of the name. The following example applies to any host within the `example.com` domain:

```
ALL : .example.com
```

- *IP address ending with a period (.)* — Placing a period at the end of an IP address matches all hosts sharing the initial numeric groups of an IP address. The following example applies to any host within the `192.168.x.x` network:

```
ALL : 192.168.
```

- *IP address/netmask pair* — Netmask expressions can also be used as a pattern to control access to a particular group of IP addresses. The following example applies to any host with an address range of `192.168.0.0` through `192.168.1.255`:

```
ALL : 192.168.0.0/255.255.254.0
```



Important

When working in the IPv4 address space, the address/prefix length (*prefixlen*) pair declarations (CIDR notation) are not supported. Only IPv6 rules can use this format.

- *[IPv6 address]/prefixlen pair* — `[net]/prefixlen` pairs can also be used as a pattern to control access to a particular group of IPv6 addresses. The following example would apply to any host with an address range of `3ffe:505:2:1::` through `3ffe:505:2:1:ffff:ffff:ffff:ffff`:

```
ALL : [3ffe:505:2:1::]/64
```

- *The asterisk (*)* — Asterisks can be used to match entire groups of hostnames or IP addresses, as long as they are not mixed in a client list containing other types of patterns. The following example would apply to any host within the `example.com` domain:

```
ALL : *.example.com
```

- *The slash (/)* — If a client list begins with a slash, it is treated as a file name. This is useful if rules specifying large numbers of hosts are necessary. The following example refers TCP Wrappers to the `/etc/telnet.hosts` file for all Telnet connections:

```
in.telnetd : /etc/telnet.hosts
```

Other, lesser used, patterns are also accepted by TCP Wrappers. Refer to the `hosts_access` man 5 page for more information.



Warning

Be very careful when using hostnames and domain names. Attackers can use a variety of tricks to circumvent accurate name resolution. In addition, disruption to DNS service prevents even authorized users from using network services. It is, therefore, best to use IP addresses whenever possible.

2.6.2.1.3. Portmap and TCP Wrappers

portmap's implementation of TCP Wrappers does not support host look-ups, which means **portmap** can not use hostnames to identify hosts. Consequently, access control rules for portmap in **hosts.allow** or **hosts.deny** must use IP addresses, or the keyword ALL, for specifying hosts.

Changes to **portmap** access control rules may not take effect immediately. You may need to restart the **portmap** service.

Widely used services, such as NIS and NFS, depend on **portmap** to operate, so be aware of these limitations.

2.6.2.1.4. Operators

At present, access control rules accept one operator, EXCEPT. It can be used in both the daemon list and the client list of a rule.

The EXCEPT operator allows specific exceptions to broader matches within the same rule.

In the following example from a **hosts.allow** file, all `example.com` hosts are allowed to connect to all services except `cracker.example.com`:

```
ALL: .example.com EXCEPT cracker.example.com
```

In another example from a **hosts.allow** file, clients from the `192.168.0.x` network can use all services except for FTP:

```
ALL EXCEPT vsftpd: 192.168.0.
```



Note

Organizationally, it is often easier to avoid using EXCEPT operators. This allows other administrators to quickly scan the appropriate files to see what hosts are allowed or denied access to services, without having to sort through EXCEPT operators.

2.6.2.2. Option Fields

In addition to basic rules that allow and deny access, the Fedora implementation of TCP Wrappers supports extensions to the access control language through *option fields*. By using option fields in hosts access rules, administrators can accomplish a variety of tasks such as altering log behavior, consolidating access control, and launching shell commands.

2.6.2.2.1. Logging

Option fields let administrators easily change the log facility and priority level for a rule by using the `severity` directive.

In the following example, connections to the SSH daemon from any host in the `example.com` domain are logged to the default `authpriv` syslog facility (because no facility value is specified) with a priority of `emerg`:

```
sshd : .example.com : severity emerg
```

It is also possible to specify a facility using the `severity` option. The following example logs any SSH connection attempts by hosts from the `example.com` domain to the `local0` facility with a priority of `alert`:

```
sshd : .example.com : severity local0.alert
```



Note

In practice, this example does not work until the syslog daemon (`syslogd`) is configured to log to the `local0` facility. Refer to the `syslog.conf` man page for information about configuring custom log facilities.

2.6.2.2.2. Access Control

Option fields also allow administrators to explicitly allow or deny hosts in a single rule by adding the `allow` or `deny` directive as the final option.

For example, the following two rules allow SSH connections from `client-1.example.com`, but deny connections from `client-2.example.com`:

```
sshd : client-1.example.com : allow
sshd : client-2.example.com : deny
```

By allowing access control on a per-rule basis, the option field allows administrators to consolidate all access rules into a single file: either `hosts.allow` or `hosts.deny`. Some administrators consider this an easier way of organizing access rules.

2.6.2.2.3. Shell Commands

Option fields allow access rules to launch shell commands through the following two directives:

- **spawn** — Launches a shell command as a child process. This directive can perform tasks like using `/usr/sbin/safe_finger` to get more information about the requesting client or create special log files using the `echo` command.

In the following example, clients attempting to access Telnet services from the `example.com` domain are quietly logged to a special file:

```
in.telnetd : .example.com \
: spawn /bin/echo `/bin/date` from %h>>/var/log/telnet.log \
```

```
: allow
```

- **twist** — Replaces the requested service with the specified command. This directive is often used to set up traps for intruders (also called "honey pots"). It can also be used to send messages to connecting clients. The **twist** directive must occur at the end of the rule line.

In the following example, clients attempting to access FTP services from the `example.com` domain are sent a message using the **echo** command:

```
vsftpd : .example.com \  
: twist /bin/echo "421 This domain has been black-listed. Access denied!"
```

For more information about shell command options, refer to the **hosts_options** man page.

2.6.2.2.4. Expansions

Expansions, when used in conjunction with the **spawn** and **twist** directives, provide information about the client, server, and processes involved.

The following is a list of supported expansions:

- **%a** — Returns the client's IP address.
- **%A** — Returns the server's IP address.
- **%c** — Returns a variety of client information, such as the username and hostname, or the username and IP address.
- **%d** — Returns the daemon process name.
- **%h** — Returns the client's hostname (or IP address, if the hostname is unavailable).
- **%H** — Returns the server's hostname (or IP address, if the hostname is unavailable).
- **%n** — Returns the client's hostname. If unavailable, `unknown` is printed. If the client's hostname and host address do not match, `paranoid` is printed.
- **%N** — Returns the server's hostname. If unavailable, `unknown` is printed. If the server's hostname and host address do not match, `paranoid` is printed.
- **%p** — Returns the daemon's process ID.
- **%s** — Returns various types of server information, such as the daemon process and the host or IP address of the server.
- **%u** — Returns the client's username. If unavailable, `unknown` is printed.

The following sample rule uses an expansion in conjunction with the **spawn** command to identify the client host in a customized log file.

When connections to the SSH daemon (`sshd`) are attempted from a host in the `example.com` domain, execute the **echo** command to log the attempt, including the client hostname (by using the **%h** expansion), to a special file:

```
sshd : .example.com \  
: twist /bin/echo "421 This domain has been black-listed. Access denied!"
```

```
: spawn /bin/echo `/bin/date` access denied to %h>>/var/log/sshd.log \
: deny
```

Similarly, expansions can be used to personalize messages back to the client. In the following example, clients attempting to access FTP services from the `example.com` domain are informed that they have been banned from the server:

```
vsftpd : .example.com \
: twist /bin/echo "421 %h has been banned from this server!"
```

For a full explanation of available expansions, as well as additional access control options, refer to section 5 of the man pages for **hosts_access** (man 5 **hosts_access**) and the man page for **hosts_options**.

Refer to [Section 2.6.5, “Additional Resources”](#) for more information about TCP Wrappers.

2.6.3. xinetd

The `xinetd` daemon is a TCP-wrapped *super service* which controls access to a subset of popular network services, including FTP, IMAP, and Telnet. It also provides service-specific configuration options for access control, enhanced logging, binding, redirection, and resource utilization control.

When a client attempts to connect to a network service controlled by `xinetd`, the super service receives the request and checks for any TCP Wrappers access control rules.

If access is allowed, `xinetd` verifies that the connection is allowed under its own access rules for that service. It also checks that the service can have more resources allotted to it and that it is not in breach of any defined rules.

If all these conditions are met (that is, access is allowed to the service; the service has not reached its resource limit; and the service is not in breach of any defined rule), `xinetd` then starts an instance of the requested service and passes control of the connection to it. After the connection has been established, `xinetd` takes no further part in the communication between the client and the server.

2.6.4. xinetd Configuration Files

The configuration files for `xinetd` are as follows:

- `/etc/xinetd.conf` — The global `xinetd` configuration file.
- `/etc/xinetd.d/` — The directory containing all service-specific files.

2.6.4.1. The `/etc/xinetd.conf` File

The `/etc/xinetd.conf` file contains general configuration settings which affect every service under `xinetd`'s control. It is read when the `xinetd` service is first started, so for configuration changes to take effect, you need to restart the `xinetd` service. The following is a sample `/etc/xinetd.conf` file:

```
defaults
{
  instances           = 60
  log_type            = SYSLOG authpriv
  log_on_success      = HOST PID
  log_on_failure      = HOST
```

```
cps                = 25 30
}
includedir /etc/xinetd.d
```

These lines control the following aspects of `xinetd`:

- `instances` — Specifies the maximum number of simultaneous requests that `xinetd` can process.
- `log_type` — Configures `xinetd` to use the **authpriv** log facility, which writes log entries to the `/var/log/secure` file. Adding a directive such as `FILE /var/log/xinetdlog` would create a custom log file called **xinetdlog** in the `/var/log/` directory.
- `log_on_success` — Configures `xinetd` to log successful connection attempts. By default, the remote host's IP address and the process ID of the server processing the request are recorded.
- `log_on_failure` — Configures `xinetd` to log failed connection attempts or if the connection was denied.
- `cps` — Configures `xinetd` to allow no more than 25 connections per second to any given service. If this limit is exceeded, the service is retired for 30 seconds.
- `includedir /etc/xinetd.d/` — Includes options declared in the service-specific configuration files located in the `/etc/xinetd.d/` directory. Refer to [Section 2.6.4.2, “The /etc/xinetd.d/ Directory”](#) for more information.

Note

Often, both the `log_on_success` and `log_on_failure` settings in `/etc/xinetd.conf` are further modified in the service-specific configuration files. More information may therefore appear in a given service's log file than the `/etc/xinetd.conf` file may indicate. Refer to [Section 2.6.4.3.1, “Logging Options”](#) for further information.

2.6.4.2. The `/etc/xinetd.d/` Directory

The `/etc/xinetd.d/` directory contains the configuration files for each service managed by `xinetd` and the names of the files correlate to the service. As with `xinetd.conf`, this directory is read only when the `xinetd` service is started. For any changes to take effect, the administrator must restart the `xinetd` service.

The format of files in the `/etc/xinetd.d/` directory use the same conventions as `/etc/xinetd.conf`. The primary reason the configuration for each service is stored in a separate file is to make customization easier and less likely to affect other services.

To gain an understanding of how these files are structured, consider the `/etc/xinetd.d/krb5-telnet` file:


```
service telnet
{
  flags                = REUSE
  socket_type         = stream
  wait                = no
  user                = root
  server              = /usr/kerberos/sbin/telnetd
```



```
log_on_failure += USERID
disable       = yes
}
```

These lines control various aspects of the **telnet** service:

- **service** — Specifies the service name, usually one of those listed in the **/etc/services** file.
- **flags** — Sets any of a number of attributes for the connection. **REUSE** instructs **xinetd** to reuse the socket for a Telnet connection.



Note
The **REUSE** flag is deprecated. All services now implicitly use the **REUSE** flag.

- **socket_type** — Sets the network socket type to **stream**.
- **wait** — Specifies whether the service is single-threaded (**yes**) or multi-threaded (**no**).
- **user** — Specifies which user ID the process runs under.
- **server** — Specifies which binary executable to launch.
- **log_on_failure** — Specifies logging parameters for **log_on_failure** in addition to those already defined in **xinetd.conf**.
- **disable** — Specifies whether the service is disabled (**yes**) or enabled (**no**).

Refer to the **xinetd.conf** man page for more information about these options and their usage.

2.6.4.3. Altering xinetd Configuration Files

A range of directives is available for services protected by **xinetd**. This section highlights some of the more commonly used options.

2.6.4.3.1. Logging Options

The following logging options are available for both **/etc/xinetd.conf** and the service-specific configuration files within the **/etc/xinetd.d/** directory.

The following is a list of some of the more commonly used logging options:

- **ATTEMPT** — Logs the fact that a failed attempt was made (**log_on_failure**).
- **DURATION** — Logs the length of time the service is used by a remote system (**log_on_success**).
- **EXIT** — Logs the exit status or termination signal of the service (**log_on_success**).
- **HOST** — Logs the remote host's IP address (**log_on_failure** and **log_on_success**).
- **PID** — Logs the process ID of the server receiving the request (**log_on_success**).
- **USERID** — Logs the remote user using the method defined in RFC 1413 for all multi-threaded stream services (**log_on_failure** and **log_on_success**).

For a complete list of logging options, refer to the `xinetd.conf` man page.

2.6.4.3.2. Access Control Options

Users of `xinetd` services can choose to use the TCP Wrappers hosts access rules, provide access control via the `xinetd` configuration files, or a mixture of both. Refer to [Section 2.6.2, “TCP Wrappers Configuration Files”](#) for more information about TCP Wrappers hosts access control files.

This section discusses using `xinetd` to control access to services.



Note

Unlike TCP Wrappers, changes to access control only take effect if the `xinetd` administrator restarts the `xinetd` service.

Also, unlike TCP Wrappers, access control through `xinetd` only affects services controlled by `xinetd`.

The `xinetd` hosts access control differs from the method used by TCP Wrappers. While TCP Wrappers places all of the access configuration within two files, `/etc/hosts.allow` and `/etc/hosts.deny`, `xinetd`'s access control is found in each service's configuration file in the `/etc/xinetd.d/` directory.

The following hosts access options are supported by `xinetd`:

- `only_from` — Allows only the specified hosts to use the service.
- `no_access` — Blocks listed hosts from using the service.
- `access_times` — Specifies the time range when a particular service may be used. The time range must be stated in 24-hour format notation, HH:MM-HH:MM.

The `only_from` and `no_access` options can use a list of IP addresses or host names, or can specify an entire network. Like TCP Wrappers, combining `xinetd` access control with the enhanced logging configuration can increase security by blocking requests from banned hosts while verbosely recording each connection attempt.

For example, the following `/etc/xinetd.d/telnet` file can be used to block Telnet access from a particular network group and restrict the overall time range that even allowed users can log in:

```
service telnet
{
  disable          = no
  flags            = REUSE
  socket_type      = stream
  wait            = no
  user             = root
  server           = /usr/kerberos/sbin/telnetd
  log_on_failure  += USERID
  no_access        = 172.16.45.0/24
  log_on_success  += PID HOST EXIT
  access_times     = 09:45-16:15
}
```

In this example, when a client system from the `10.0.1.0/24` network, such as `10.0.1.2`, tries to access the Telnet service, it receives the following message:

Connection closed by foreign host.

In addition, their login attempts are logged in `/var/log/messages` as follows:

```
Sep  7 14:58:33 localhost xinetd[5285]: FAIL: telnet address from=172.16.45.107
Sep  7 14:58:33 localhost xinetd[5283]: START: telnet pid=5285 from=172.16.45.107
Sep  7 14:58:33 localhost xinetd[5283]: EXIT: telnet status=0 pid=5285 duration=0(sec)
```

When using TCP Wrappers in conjunction with `xinetd` access controls, it is important to understand the relationship between the two access control mechanisms.

The following is the sequence of events followed by `xinetd` when a client requests a connection:

1. The `xinetd` daemon accesses the TCP Wrappers hosts access rules using a `libwrap.a` library call. If a deny rule matches the client, the connection is dropped. If an allow rule matches the client, the connection is passed to `xinetd`.
2. The `xinetd` daemon checks its own access control rules both for the `xinetd` service and the requested service. If a deny rule matches the client, the connection is dropped. Otherwise, `xinetd` starts an instance of the requested service and passes control of the connection to that service.



Important

Care should be taken when using TCP Wrappers access controls in conjunction with `xinetd` access controls. Misconfiguration can cause undesirable effects.

2.6.4.3.3. Binding and Redirection Options

The service configuration files for `xinetd` support binding the service to an IP address and redirecting incoming requests for that service to another IP address, hostname, or port.

Binding is controlled with the `bind` option in the service-specific configuration files and links the service to one IP address on the system. When this is configured, the `bind` option only allows requests to the correct IP address to access the service. You can use this method to bind different services to different network interfaces based on requirements.

This is particularly useful for systems with multiple network adapters or with multiple IP addresses. On such a system, insecure services (for example, Telnet), can be configured to listen only on the interface connected to a private network and not to the interface connected to the Internet.

The `redirect` option accepts an IP address or hostname followed by a port number. It configures the service to redirect any requests for this service to the specified host and port number. This feature can be used to point to another port number on the same system, redirect the request to a different IP address on the same machine, shift the request to a totally different system and port number, or any combination of these options. A user connecting to a certain service on a system may therefore be rerouted to another system without disruption.

The `xinetd` daemon is able to accomplish this redirection by spawning a process that stays alive for the duration of the connection between the requesting client machine and the host actually providing the service, transferring data between the two systems.

The advantages of the `bind` and `redirect` options are most clearly evident when they are used together. By binding a service to a particular IP address on a system and then redirecting requests for this service to a second machine that only the first machine can see, an internal system can be used to provide services for a totally different network. Alternatively, these options can be used to limit the exposure of a particular service on a multi-homed machine to a known IP address, as well as redirect any requests for that service to another machine especially configured for that purpose.

For example, consider a system that is used as a firewall with this setting for its Telnet service:

```
service telnet
{
  socket_type = stream
  wait       = no
  server     = /usr/kerberos/sbin/telnetd
  log_on_success += DURATION USERID
  log_on_failure += USERID
  bind       = 123.123.123.123
  redirect   = 10.0.1.13 23
}
```

The `bind` and `redirect` options in this file ensure that the Telnet service on the machine is bound to the external IP address (123.123.123.123), the one facing the Internet. In addition, any requests for Telnet service sent to 123.123.123.123 are redirected via a second network adapter to an internal IP address (10.0.1.13) that only the firewall and internal systems can access. The firewall then sends the communication between the two systems, and the connecting system thinks it is connected to 123.123.123.123 when it is actually connected to a different machine.

This feature is particularly useful for users with broadband connections and only one fixed IP address. When using Network Address Translation (NAT), the systems behind the gateway machine, which are using internal-only IP addresses, are not available from outside the gateway system. However, when certain services controlled by `xinetd` are configured with the `bind` and `redirect` options, the gateway machine can act as a proxy between outside systems and a particular internal machine configured to provide the service. In addition, the various `xinetd` access control and logging options are also available for additional protection.

2.6.4.3.4. Resource Management Options

The `xinetd` daemon can add a basic level of protection from Denial of Service (DoS) attacks. The following is a list of directives which can aid in limiting the effectiveness of such attacks:

- `per_source` — Defines the maximum number of instances for a service per source IP address. It accepts only integers as an argument and can be used in both `xinetd.conf` and in the service-specific configuration files in the `xinetd.d/` directory.
- `cps` — Defines the maximum number of connections per second. This directive takes two integer arguments separated by white space. The first argument is the maximum number of connections allowed to the service per second. The second argument is the number of seconds that `xinetd` must wait before re-enabling the service. It accepts only integers as arguments and can be used in either the `xinetd.conf` file or the service-specific configuration files in the `xinetd.d/` directory.
- `max_load` — Defines the CPU usage or load average threshold for a service. It accepts a floating point number argument.

The load average is a rough measure of how many processes are active at a given time. See the `uptime`, `who`, and `procinfo` commands for more information about load average.

There are more resource management options available for `xinetd`. Refer to the `xinetd.conf` man page for more information.

2.6.5. Additional Resources

More information about TCP Wrappers and `xinetd` is available from system documentation and on the Internet.

2.6.5.1. Installed TCP Wrappers Documentation

The documentation on your system is a good place to start looking for additional configuration options for TCP Wrappers, `xinetd`, and access control.

- `/usr/share/doc/tcp_wrappers-<version>/` — This directory contains a **README** file that discusses how TCP Wrappers work and the various hostname and host address spoofing risks that exist.
- `/usr/share/doc/xinetd-<version>/` — This directory contains a **README** file that discusses aspects of access control and a **sample.conf** file with various ideas for modifying service-specific configuration files in the `/etc/xinetd.d/` directory.
- TCP Wrappers and `xinetd`-related man pages — A number of man pages exist for the various applications and configuration files involved with TCP Wrappers and `xinetd`. The following are some of the more important man pages:

Server Applications

- `man xinetd` — The man page for `xinetd`.

Configuration Files

- `man 5 hosts_access` — The man page for the TCP Wrappers hosts access control files.
- `man hosts_options` — The man page for the TCP Wrappers options fields.
- `man xinetd.conf` — The man page listing `xinetd` configuration options.

2.6.5.2. Useful TCP Wrappers Websites

- <http://www.xinetd.org>^A — The home of `xinetd`, containing sample configuration files, a full listing of features, and an informative FAQ.
- <http://www.docstoc.com/docs/2133633/An-Unofficial-Xinetd-Tutorial> — A thorough tutorial that discusses many different ways to optimize default `xinetd` configuration files to meet specific security goals.

2.6.5.3. Related Books

- *Hacking Linux Exposed* by Brian Hatch, James Lee, and George Kurtz; Osbourne/McGraw-Hill — An excellent security resource with information about TCP Wrappers and `xinetd`.

2.7. Kerberos

System security and integrity within a network can be unwieldy. It can occupy the time of several administrators just to keep track of what services are being run on a network and the manner in which these services are used.

Further, authenticating users to network services can prove dangerous when the method used by the protocol is inherently insecure, as evidenced by the transfer of unencrypted passwords over a network using the traditional FTP and Telnet protocols.

Kerberos is a way to eliminate the need for protocols that allow unsafe methods of authentication, thereby enhancing overall network security.

2.7.1. What is Kerberos?

Kerberos is a network authentication protocol created by MIT, and uses symmetric-key cryptography⁵ to authenticate users to network services, which means passwords are never actually sent over the network.

Consequently, when users authenticate to network services using Kerberos, unauthorized users attempting to gather passwords by monitoring network traffic are effectively thwarted.

2.7.1.1. Advantages of Kerberos

Most conventional network services use password-based authentication schemes. Such schemes require a user to authenticate to a given network server by supplying their username and password. Unfortunately, the transmission of authentication information for many services is unencrypted. For such a scheme to be secure, the network has to be inaccessible to outsiders, and all computers and users on the network must be trusted and trustworthy.

Even if this is the case, a network that is connected to the Internet can no longer be assumed to be secure. Any attacker who gains access to the network can use a simple packet analyzer, also known as a packet sniffer, to intercept usernames and passwords, compromising user accounts and the integrity of the entire security infrastructure.

The primary design goal of Kerberos is to eliminate the transmission of unencrypted passwords across the network. If used properly, Kerberos effectively eliminates the threat that packet sniffers would otherwise pose on a network.

2.7.1.2. Disadvantages of Kerberos

Although Kerberos removes a common and severe security threat, it may be difficult to implement for a variety of reasons:

- Migrating user passwords from a standard UNIX password database, such as `/etc/passwd` or `/etc/shadow`, to a Kerberos password database can be tedious, as there is no automated mechanism to perform this task. Refer to Question 2.23 in the online Kerberos FAQ:
<http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html>⁶
- Kerberos has only partial compatibility with the Pluggable Authentication Modules (PAM) system used by most Fedora servers. Refer to [Section 2.7.4, “Kerberos and PAM”](#) for more information about this issue.
- Kerberos assumes that each user is trusted but is using an untrusted host on an untrusted network. Its primary goal is to prevent unencrypted passwords from being transmitted across that network. However, if anyone other than the proper user has access to the one host that issues tickets used

⁵ A system where both the client and the server share a common key that is used to encrypt and decrypt network communication.

for authentication — called the *key distribution center (KDC)* — the entire Kerberos authentication system is at risk.

- For an application to use Kerberos, its source must be modified to make the appropriate calls into the Kerberos libraries. Applications modified in this way are considered to be *Kerberos-aware*, or *kerberized*. For some applications, this can be quite problematic due to the size of the application or its design. For other incompatible applications, changes must be made to the way in which the server and client communicate. Again, this may require extensive programming. Closed-source applications that do not have Kerberos support by default are often the most problematic.
- Kerberos is an all-or-nothing solution. If Kerberos is used on the network, any unencrypted passwords transferred to a non-Kerberos aware service is at risk. Thus, the network gains no benefit from the use of Kerberos. To secure a network with Kerberos, one must either use Kerberos-aware versions of *all* client/server applications that transmit passwords unencrypted, or not use *any* such client/server applications at all.

2.7.2. Kerberos Terminology

Kerberos has its own terminology to define various aspects of the service. Before learning how Kerberos works, it is important to learn the following terms.

authentication server (AS)

A server that issues tickets for a desired service which are in turn given to users for access to the service. The AS responds to requests from clients who do not have or do not send credentials with a request. It is usually used to gain access to the ticket-granting server (TGS) service by issuing a ticket-granting ticket (TGT). The AS usually runs on the same host as the key distribution center (KDC).

ciphertext

Encrypted data.

client

An entity on the network (a user, a host, or an application) that can receive a ticket from Kerberos.

credentials

A temporary set of electronic credentials that verify the identity of a client for a particular service. Also called a ticket.

credential cache or ticket file

A file which contains the keys for encrypting communications between a user and various network services. Kerberos 5 supports a framework for using other cache types, such as shared memory, but files are more thoroughly supported.

crypt hash

A one-way hash used to authenticate users. These are more secure than using unencrypted data, but they are still relatively easy to decrypt for an experienced cracker.

GSS-API

The Generic Security Service Application Program Interface (defined in RFC-2743 published by The Internet Engineering Task Force) is a set of functions which provide security services. This API is used by clients and services to authenticate to each other without either program having specific knowledge of the underlying mechanism. If a network service (such as cyrus-IMAP) uses GSS-API, it can authenticate using Kerberos.

hash

Also known as a *hash value*. A value generated by passing a string through a *hash function*. These values are typically used to ensure that transmitted data has not been tampered with.

hash function

A way of generating a digital "fingerprint" from input data. These functions rearrange, transpose or otherwise alter data to produce a *hash value*.

key

Data used when encrypting or decrypting other data. Encrypted data cannot be decrypted without the proper key or extremely good fortune on the part of the cracker.

key distribution center (KDC)

A service that issues Kerberos tickets, and which usually run on the same host as the ticket-granting server (TGS).

keytab (or key table)

A file that includes an unencrypted list of principals and their keys. Servers retrieve the keys they need from keytab files instead of using **kinit**. The default keytab file is **/etc/krb5.keytab**. The KDC administration server, **/usr/kerberos/sbin/kadmind**, is the only service that uses any other file (it uses **/var/kerberos/krb5kdc/kadm5.keytab**).

kinit

The **kinit** command allows a principal who has already logged in to obtain and cache the initial ticket-granting ticket (TGT). Refer to the **kinit** man page for more information.

principal (or principal name)

The principal is the unique name of a user or service allowed to authenticate using Kerberos. A principal follows the form `root[/instance]@REALM`. For a typical user, the root is the same as their login ID. The `instance` is optional. If the principal has an instance, it is separated from the root with a forward slash ("/"). An empty string ("") is considered a valid instance (which differs from the default NULL instance), but using it can be confusing. All principals in a realm have their own key, which for users is derived from a password or is randomly set for services.

realm

A network that uses Kerberos, composed of one or more servers called KDCs and a potentially large number of clients.

service

A program accessed over the network.

ticket

A temporary set of electronic credentials that verify the identity of a client for a particular service. Also called credentials.

ticket-granting server (TGS)

A server that issues tickets for a desired service which are in turn given to users for access to the service. The TGS usually runs on the same host as the KDC.

ticket-granting ticket (TGT)

A special ticket that allows the client to obtain additional tickets without applying for them from the KDC.

unencrypted password

A plain text, human-readable password.

2.7.3. How Kerberos Works

Kerberos differs from username/password authentication methods. Instead of authenticating each user to each network service, Kerberos uses symmetric encryption and a trusted third party (a KDC), to authenticate users to a suite of network services. When a user authenticates to the KDC, the KDC sends a ticket specific to that session back to the user's machine, and any Kerberos-aware services look for the ticket on the user's machine rather than requiring the user to authenticate using a password.

When a user on a Kerberos-aware network logs in to their workstation, their principal is sent to the KDC as part of a request for a TGT from the Authentication Server. This request can be sent by the log-in program so that it is transparent to the user, or can be sent by the **kinit** program after the user logs in.

The KDC then checks for the principal in its database. If the principal is found, the KDC creates a TGT, which is encrypted using the user's key and returned to that user.

The login or **kinit** program on the client then decrypts the TGT using the user's key, which it computes from the user's password. The user's key is used only on the client machine and is *not* transmitted over the network.

The TGT is set to expire after a certain period of time (usually ten to twenty-four hours) and is stored in the client machine's credentials cache. An expiration time is set so that a compromised TGT is of use to an attacker for only a short period of time. After the TGT has been issued, the user does not have to re-enter their password until the TGT expires or until they log out and log in again.

Whenever the user needs access to a network service, the client software uses the TGT to request a new ticket for that specific service from the TGS. The service ticket is then used to authenticate the user to that service transparently.



Warning

The Kerberos system can be compromised if a user on the network authenticates against a non-Kerberos aware service by transmitting a password in plain text. The use of non-Kerberos aware services is highly discouraged. Such services include Telnet and FTP. The use of other encrypted protocols, such as SSH or SSL-secured services, however, is preferred, although not ideal.

This is only a broad overview of how Kerberos authentication works. Refer to [Section 2.7.10, "Additional Resources"](#) for links to more in-depth information.



Note

Kerberos depends on the following network services to function correctly.

- Approximate clock synchronization between the machines on the network.

A clock synchronization program should be set up for the network, such as **ntpd**. Refer to `/usr/share/doc/ntp-<version-number>/index.html` for details on setting up Network Time Protocol servers (where `<version-number>` is the version number of the **ntp** package installed on your system).

- Domain Name Service (DNS).

You should ensure that the DNS entries and hosts on the network are all properly configured. Refer to the *Kerberos V5 System Administrator's Guide* in `/usr/share/doc/krb5-server-<version-number>` for more information (where `<version-number>` is the version number of the `krb5-server` package installed on your system).

2.7.4. Kerberos and PAM

Kerberos-aware services do not currently make use of Pluggable Authentication Modules (PAM) — these services bypass PAM completely. However, applications that use PAM can make use of Kerberos for authentication if the `pam_krb5` module (provided in the `pam_krb5` package) is installed. The `pam_krb5` package contains sample configuration files that allow services such as `login` and `gdm` to authenticate users as well as obtain initial credentials using their passwords. If access to network servers is always performed using Kerberos-aware services or services that use GSS-API, such as IMAP, the network can be considered reasonably safe.



Important

Administrators should be careful not to allow users to authenticate to most network services using Kerberos passwords. Many protocols used by these services do not encrypt the password before sending it over the network, destroying the benefits of the Kerberos system. For example, users should not be allowed to authenticate to Telnet services with the same password they use for Kerberos authentication.

2.7.5. Configuring a Kerberos 5 Server

When setting up Kerberos, install the KDC first. If it is necessary to set up slave servers, install the master first.

To configure the first Kerberos KDC, follow these steps:

1. Ensure that time synchronization and DNS are functioning correctly on all client and server machines before configuring Kerberos. Pay particular attention to time synchronization between the Kerberos server and its clients. If the time difference between the server and client is greater than five minutes (this is configurable in Kerberos 5), Kerberos clients can not authenticate to the server. This time synchronization is necessary to prevent an attacker from using an old Kerberos ticket to masquerade as a valid user.

It is advisable to set up a Network Time Protocol (NTP) compatible client/server network even if Kerberos is not being used. Fedora includes the `ntp` package for this purpose. Refer to `/usr/share/doc/ntp-<version-number>/index.html` (where `<version-number>` is the version number of the `ntp` package installed on your system) for details about how to set up Network Time Protocol servers, and <http://www.ntp.org> for more information about NTP.

2. Install the `krb5-libs`, `krb5-server`, and `krb5-workstation` packages on the dedicated machine which runs the KDC. This machine needs to be very secure — if possible, it should not run any services other than the KDC.

3. Edit the `/etc/krb5.conf` and `/var/kerberos/krb5kdc/kdc.conf` configuration files to reflect the realm name and domain-to-realm mappings. A simple realm can be constructed by replacing instances of `EXAMPLE.COM` and `example.com` with the correct domain name — being certain to keep uppercase and lowercase names in the correct format — and by changing the KDC from `kerberos.example.com` to the name of the Kerberos server. By convention, all realm names are uppercase and all DNS hostnames and domain names are lowercase. For full details about the formats of these configuration files, refer to their respective man pages.
4. Create the database using the `kdb5_util` utility from a shell prompt:

```
/usr/kerberos/sbin/kdb5_util create -s
```

The `create` command creates the database that stores keys for the Kerberos realm. The `-s` switch forces creation of a `stash` file in which the master server key is stored. If no `stash` file is present from which to read the key, the Kerberos server (`krb5kdc`) prompts the user for the master server password (which can be used to regenerate the key) every time it starts.

5. Edit the `/var/kerberos/krb5kdc/kadm5.ac1` file. This file is used by `kadmin` to determine which principals have administrative access to the Kerberos database and their level of access. Most organizations can get by with a single line:

```
*/admin@EXAMPLE.COM *
```

Most users are represented in the database by a single principal (with a `NULL`, or empty, instance, such as `joe@EXAMPLE.COM`). In this configuration, users with a second principal with an instance of `admin` (for example, `joe/admin@EXAMPLE.COM`) are able to wield full power over the realm's Kerberos database.

After `kadmin` has been started on the server, any user can access its services by running `kadmin` on any of the clients or servers in the realm. However, only users listed in the `kadm5.ac1` file can modify the database in any way, except for changing their own passwords.



Note

The `kadmin` utility communicates with the `kadmin` server over the network, and uses Kerberos to handle authentication. Consequently, the first principal must already exist before connecting to the server over the network to administer it. Create the first principal with the `kadmin.local` command, which is specifically designed to be used on the same host as the KDC and does not use Kerberos for authentication.

Type the following `kadmin.local` command at the KDC terminal to create the first principal:

```
/usr/kerberos/sbin/kadmin.local -q "addprinc username/admin"
```

6. Start Kerberos using the following commands:

```
/sbin/service krb5kdc start
/sbin/service kadmin start
/sbin/service krb524 start
```

7. Add principals for the users using the **addprinc** command within **kadmin**. **kadmin** and **kadmin.local** are command line interfaces to the KDC. As such, many commands — such as **addprinc** — are available after launching the **kadmin** program. Refer to the **kadmin** man page for more information.
8. Verify that the KDC is issuing tickets. First, run **kinit** to obtain a ticket and store it in a credential cache file. Next, use **klist** to view the list of credentials in the cache and use **kdestroy** to destroy the cache and the credentials it contains.



Note

By default, **kinit** attempts to authenticate using the same system login username (not the Kerberos server). If that username does not correspond to a principal in the Kerberos database, **kinit** issues an error message. If that happens, supply **kinit** with the name of the correct principal as an argument on the command line (**kinit <principal>**).

Once these steps are completed, the Kerberos server should be up and running.

2.7.6. Configuring a Kerberos 5 Client

Setting up a Kerberos 5 client is less involved than setting up a server. At a minimum, install the client packages and provide each client with a valid **krb5.conf** configuration file. While **ssh** and **slogin** are the preferred method of remotely logging in to client systems, Kerberized versions of **rsh** and **rlogin** are still available, though deploying them requires that a few more configuration changes be made.

1. Be sure that time synchronization is in place between the Kerberos client and the KDC. Refer to [Section 2.7.5, “Configuring a Kerberos 5 Server”](#) for more information. In addition, verify that DNS is working properly on the Kerberos client before configuring the Kerberos client programs.
2. Install the **krb5-libs** and **krb5-workstation** packages on all of the client machines. Supply a valid **/etc/krb5.conf** file for each client (usually this can be the same **krb5.conf** file used by the KDC).
3. Before a workstation in the realm can use Kerberos to authenticate users who connect using **ssh** or Kerberized **rsh** or **rlogin**, it must have its own host principal in the Kerberos database. The **sshd**, **kshd**, and **klogind** server programs all need access to the keys for the *host* service's principal. Additionally, in order to use the kerberized **rsh** and **rlogin** services, that workstation must have the **xinetd** package installed.

Using **kadmin**, add a host principal for the workstation on the KDC. The instance in this case is the hostname of the workstation. Use the **-randkey** option for the **kadmin**'s **addprinc** command to create the principal and assign it a random key:

```
addprinc -randkey host/blah.example.com
```

Now that the principal has been created, keys can be extracted for the workstation by running **kadmin** on the workstation itself, and using the **ktadd** command within **kadmin**:

```
ktadd -k /etc/krb5.keytab host/blah.example.com
```

4. To use other kerberized network services, they must first be started. Below is a list of some common kerberized services and instructions about enabling them:
- **ssh** — OpenSSH uses GSS-API to authenticate users to servers if the client's and server's configuration both have `GSSAPIAuthentication` enabled. If the client also has `GSSAPIDelegateCredentials` enabled, the user's credentials are made available on the remote system.
 - **rsh** and **rlogin** — To use the kerberized versions of **rsh** and **rlogin**, enable **klogin**, **eklogin**, and **kshell**.
 - Telnet — To use kerberized Telnet, **krb5-telnet** must be enabled.
 - FTP — To provide FTP access, create and extract a key for the principal with a root of `ftp`. Be certain to set the instance to the fully qualified hostname of the FTP server, then enable **gssftp**.
 - IMAP — To use a kerberized IMAP server, the **cyrus-imap** package uses Kerberos 5 if it also has the **cyrus-sasl-gssapi** package installed. The **cyrus-sasl-gssapi** package contains the Cyrus SASL plugins which support GSS-API authentication. Cyrus IMAP should function properly with Kerberos as long as the **cyrus** user is able to find the proper key in `/etc/krb5.keytab`, and the root for the principal is set to **imap** (created with **kadmin**).
- An alternative to **cyrus-imap** can be found in the **dovecot** package, which is also included in Fedora. This package contains an IMAP server but does not, to date, support GSS-API and Kerberos.
- CVS — To use a kerberized CVS server, **gserver** uses a principal with a root of `cv`s and is otherwise identical to the CVS **pserver**.

2.7.7. Domain-to-Realm Mapping

When a client attempts to access a service running on a particular server, it knows the name of the service (*host*) and the name of the server (*foo.example.com*), but because more than one realm may be deployed on your network, it must guess at the name of the realm in which the service resides.

By default, the name of the realm is taken to be the DNS domain name of the server, upper-cased.

```
foo.example.org → EXAMPLE.ORG
foo.example.com → EXAMPLE.COM
foo.hq.example.com → HQ.EXAMPLE.COM
```

In some configurations, this will be sufficient, but in others, the realm name which is derived will be the name of a non-existent realm. In these cases, the mapping from the server's DNS domain name to the name of its realm must be specified in the *domain_realm* section of the client system's **krb5.conf**. For example:

```
[domain_realm]
.example.com = EXAMPLE.COM
example.com = EXAMPLE.COM
```

The above configuration specifies two mappings. The first mapping specifies that any system in the "example.com" DNS domain belongs to the *EXAMPLE.COM* realm. The second specifies that a system with the exact name "example.com" is also in the realm. (The distinction between a domain

and a specific host is marked by the presence or lack of an initial ".") The mapping can also be stored directly in DNS.

2.7.8. Setting Up Secondary KDCs

For a number of reasons, you may choose to run multiple KDCs for a given realm. In this scenario, one KDC (the *master KDC*) keeps a writable copy of the realm database and runs **kadmind** (it is also your realm's *admin server*), and one or more KDCs (*slave KDCs*) keep read-only copies of the database and run **kpropd**.

The master-slave propagation procedure entails the master KDC dumping its database to a temporary dump file and then transmitting that file to each of its slaves, which then overwrite their previously-received read-only copies of the database with the contents of the dump file.

To set up a slave KDC, first ensure that the master KDC's **krb5.conf** and **kdc.conf** files are copied to the slave KDC.

Start **kadmin.local** from a root shell on the master KDC and use its **add_principal** command to create a new entry for the master KDC's *host* service, and then use its **ktadd** command to simultaneously set a random key for the service and store the random key in the master's default keytab file. This key will be used by the **kprop** command to authenticate to the slave servers. You will only need to do this once, regardless of how many slave servers you install.

```
# kadmin.local -r EXAMPLE.COM

Authenticating as principal root/admin@EXAMPLE.COM with password.

kadmin: add_principal -randkey host/masterkdc.example.com

Principal "host/host/masterkdc.example.com@EXAMPLE.COM" created.

kadmin: ktadd host/masterkdc.example.com

Entry for principal host/masterkdc.example.com with kvno 3, encryption type Triple DES cbc
mode with HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/masterkdc.example.com with kvno 3, encryption type ArcFour with HMAC/
md5 added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/masterkdc.example.com with kvno 3, encryption type DES with HMAC/sha1
added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/masterkdc.example.com with kvno 3, encryption type DES cbc mode with
RSA-MD5 added to keytab WRFILE:/etc/krb5.keytab.

kadmin: quit
```

Start **kadmin** from a root shell on the slave KDC and use its **add_principal** command to create a new entry for the slave KDC's *host* service, and then use **kadmin's ktadd** command to simultaneously set a random key for the service and store the random key in the slave's default keytab file. This key is used by the **kpropd** service when authenticating clients.

```
# kadmin -p jimbo/admin@EXAMPLE.COM -r EXAMPLE.COM

Authenticating as principal jimbo/admin@EXAMPLE.COM with password.

Password for jimbo/admin@EXAMPLE.COM:
```

```

kadmin: add_principal -randkey host/slavekdc.example.com

Principal "host/slavekdc.example.com@EXAMPLE.COM" created.

kadmin: ktadd host/slavekdc.example.com@EXAMPLE.COM

Entry for principal host/slavekdc.example.com with kvno 3, encryption type Triple DES cbc mode
with HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/slavekdc.example.com with kvno 3, encryption type ArcFour with HMAC/
md5 added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/slavekdc.example.com with kvno 3, encryption type DES with HMAC/sha1
added to keytab WRFILE:/etc/krb5.keytab.

Entry for principal host/slavekdc.example.com with kvno 3, encryption type DES cbc mode with
RSA-MD5 added to keytab WRFILE:/etc/krb5.keytab.

kadmin: quit

```

With its service key, the slave KDC could authenticate any client which would connect to it. Obviously, not all of them should be allowed to provide the slave's **kprop** service with a new realm database. To restrict access, the **kprop** service on the slave KDC will only accept updates from clients whose principal names are listed in `/var/kerberos/krb5kdc/kpropd.ac1`. Add the master KDC's host service's name to that file.

```
# echo host/masterkdc.example.com@EXAMPLE.COM > /var/kerberos/krb5kdc/kpropd.ac1
```

Once the slave KDC has obtained a copy of the database, it will also need the master key which was used to encrypt it. If your KDC database's master key is stored in a *stash* file on the master KDC (typically named `/var/kerberos/krb5kdc/.k5.REALM`, either copy it to the slave KDC using any available secure method, or create a dummy database and identical stash file on the slave KDC by running **kdb5_util create -s** (the dummy database will be overwritten by the first successful database propagation) and supplying the same password.

Ensure that the slave KDC's firewall allows the master KDC to contact it using TCP on port 754 (*krb5_prop*), and start the **kprop** service. Then, double-check that the **kadmin** service is *disabled*.

Now perform a manual database propagation test by dumping the realm database, on the master KDC, to the default data file which the **kprop** command will read (`/var/kerberos/krb5kdc/slave_datatrans`), and then use the **kprop** command to transmit its contents to the slave KDC.

```
# /usr/kerberos/sbin/kdb5_util dump /var/kerberos/krb5kdc/slave_datatrans# kprop
slavekdc.example.com
```

Using **kinit**, verify that a client system whose **krb5.conf** lists only the slave KDC in its list of KDCs for your realm is now correctly able to obtain initial credentials from the slave KDC.

That done, simply create a script which dumps the realm database and runs the **kprop** command to transmit the database to each slave KDC in turn, and configure the **cron** service to run the script periodically.

2.7.9. Setting Up Cross Realm Authentication

Cross-realm authentication is the term which is used to describe situations in which clients (typically users) of one realm use Kerberos to authenticate to services (typically server processes running on a particular server system) which belong to a realm other than their own.

For the simplest case, in order for a client of a realm named **A.EXAMPLE.COM** to access a service in the **B.EXAMPLE.COM** realm, both realms must share a key for a principal named **krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM**, and both keys must have the same key version number associated with them.

To accomplish this, select a very strong password or passphrase, and create an entry for the principal in both realms using `kadmin`.

```
# kadmin -r A.EXAMPLE.COM kadmin: add_principal krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM Enter
password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM": Re-enter password for principal
"krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM": Principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM" created.
quit # kadmin -r B.EXAMPLE.COM kadmin: add_principal krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM Enter
password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM": Re-enter password for principal
"krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM": Principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM" created.
quit
```

Use the `get_principal` command to verify that both entries have matching key version numbers (`kvno` values) and encryption types.



Dumping the Database Doesn't Do It

Security-conscious administrators may attempt to use the `add_principal` command's `-randkey` option to assign a random key instead of a password, dump the new entry from the database of the first realm, and import it into the second. This will not work unless the master keys for the realm databases are identical, as the keys contained in a database dump are themselves encrypted using the master key.

Clients in the **A.EXAMPLE.COM** realm are now able to authenticate to services in the **B.EXAMPLE.COM** realm. Put another way, the **B.EXAMPLE.COM** realm now *trusts* the **A.EXAMPLE.COM** realm, or phrased even more simply, **B.EXAMPLE.COM** now *trusts* **A.EXAMPLE.COM**.

This brings us to an important point: cross-realm trust is unidirectional by default. The KDC for the **B.EXAMPLE.COM** realm may trust clients from the **A.EXAMPLE.COM** to authenticate to services in the **B.EXAMPLE.COM** realm, but the fact that it does has no effect on whether or not clients in the **B.EXAMPLE.COM** realm are trusted to authenticate to services in the **A.EXAMPLE.COM** realm. To establish trust in the other direction, both realms would need to share keys for the **krbtgt/A.EXAMPLE.COM@B.EXAMPLE.COM** service (take note of the reversed in order of the two realms compared to the example above).

If direct trust relationships were the only method for providing trust between realms, networks which contain multiple realms would be very difficult to set up. Luckily, cross-realm trust is transitive. If clients from **A.EXAMPLE.COM** can authenticate to services in **B.EXAMPLE.COM**, and clients from **B.EXAMPLE.COM** can authenticate to services in **C.EXAMPLE.COM**, then clients in **A.EXAMPLE.COM** can also authenticate to services in **C.EXAMPLE.COM**, *even if C.EXAMPLE.COM doesn't directly trust A.EXAMPLE.COM*. This means that, on a network with multiple realms which all need to trust each other, making good choices about which trust relationships to set up can greatly reduce the amount of effort required.

Now you face the more conventional problems: the client's system must be configured so that it can properly deduce the realm to which a particular service belongs, and it must be able to determine how to obtain credentials for services in that realm.

First things first: the principal name for a service provided from a specific server system in a given realm typically looks like this:


```
service/server.example.com@EXAMPLE.COM
```

In this example, *service* is typically either the name of the protocol in use (other common values include *ldap*, *imap*, *cvs*, and *HTTP*) or *host*, *server.example.com* is the fully-qualified domain name of the system which runs the service, and **EXAMPLE.COM** is the name of the realm.

To deduce the realm to which the service belongs, clients will most often consult DNS or the **domain_realm** section of **/etc/krb5.conf** to map either a hostname (*server.example.com*) or a DNS domain name (*.example.com*) to the name of a realm (*EXAMPLE.COM*).

Having determined which to which realm a service belongs, a client then has to determine the set of realms which it needs to contact, and in which order it must contact them, to obtain credentials for use in authenticating to the service.

This can be done in one of two ways.

The default method, which requires no explicit configuration, is to give the realms names within a shared hierarchy. For an example, assume realms named **A.EXAMPLE.COM**, **B.EXAMPLE.COM**, and **EXAMPLE.COM**. When a client in the **A.EXAMPLE.COM** realm attempts to authenticate to a service in **B.EXAMPLE.COM**, it will, by default, first attempt to get credentials for the **EXAMPLE.COM** realm, and then to use those credentials to obtain credentials for use in the **B.EXAMPLE.COM** realm.

The client in this scenario treats the realm name as one might treat a DNS name. It repeatedly strips off the components of its own realm's name to generate the names of realms which are "above" it in the hierarchy until it reaches a point which is also "above" the service's realm. At that point it begins prepending components of the service's realm name until it reaches the service's realm. Each realm which is involved in the process is another "hop".

For example, using credentials in **A.EXAMPLE.COM**, authenticating to a service in

B.EXAMPLE.COM
A.EXAMPLE.COM → **EXAMPLE.COM** → **B.EXAMPLE.COM**

- **A.EXAMPLE.COM** and **EXAMPLE.COM** share a key for **krbtgt/EXAMPLE.COM@A.EXAMPLE.COM**
- **EXAMPLE.COM** and **B.EXAMPLE.COM** share a key for **krbtgt/B.EXAMPLE.COM@EXAMPLE.COM**

Another example, using credentials in **SITE1.SALES.EXAMPLE.COM**, authenticating to a service in

EVERYWHERE.EXAMPLE.COM
SITE1.SALES.EXAMPLE.COM → **SALES.EXAMPLE.COM** → **EXAMPLE.COM** → **EVERYWHERE.EXAMPLE.COM**

- **SITE1.SALES.EXAMPLE.COM** and **SALES.EXAMPLE.COM** share a key for **krbtgt/SALES.EXAMPLE.COM@SITE1.SALES.EXAMPLE.COM**
- **SALES.EXAMPLE.COM** and **EXAMPLE.COM** share a key for **krbtgt/EXAMPLE.COM@SALES.EXAMPLE.COM**
- **EXAMPLE.COM** and **EVERYWHERE.EXAMPLE.COM** share a key for **krbtgt/EVERYWHERE.EXAMPLE.COM@EXAMPLE.COM**

Another example, this time using realm names whose names share no common suffix

(**DEVEL.EXAMPLE.COM** and **PROD.EXAMPLE.ORG**)
DEVEL.EXAMPLE.COM → **EXAMPLE.COM** → **COM** → **ORG** → **EXAMPLE.ORG** → **PROD.EXAMPLE.ORG**

- **DEVEL.EXAMPLE.COM** and **EXAMPLE.COM** share a key for **krbtgt/EXAMPLE.COM@DEVEL.EXAMPLE.COM**
- **EXAMPLE.COM** and **COM** share a key for **krbtgt/COM@EXAMPLE.COM**

- **COM** and **ORG** share a key for **krbtgt/ORG@COM**
- **ORG** and **EXAMPLE.ORG** share a key for **krbtgt/EXAMPLE.ORG@ORG**
- **EXAMPLE.ORG** and **PROD.EXAMPLE.ORG** share a key for **krbtgt/PROD.EXAMPLE.ORG@EXAMPLE.ORG**

The more complicated, but also more flexible, method involves configuring the **capaths** section of **/etc/krb5.conf**, so that clients which have credentials for one realm will be able to look up which realm is next in the chain which will eventually lead to the being able to authenticate to servers.

The format of the **capaths** section is relatively straightforward: each entry in the section is named after a realm in which a client might exist. Inside of that subsection, the set of intermediate realms from which the client must obtain credentials is listed as values of the key which corresponds to the realm in which a service might reside. If there are no intermediate realms, the value "." is used.

Here's an example:

```
[capaths]
A.EXAMPLE.COM = {
B.EXAMPLE.COM = .
C.EXAMPLE.COM = B.EXAMPLE.COM
D.EXAMPLE.COM = B.EXAMPLE.COM
D.EXAMPLE.COM = C.EXAMPLE.COM
}
```

In this example, clients in the **A.EXAMPLE.COM** realm can obtain cross-realm credentials for **B.EXAMPLE.COM** directly from the **A.EXAMPLE.COM** KDC.

If those clients wish to contact a service in the **C.EXAMPLE.COM** realm, they will first need to obtain necessary credentials from the **B.EXAMPLE.COM** realm (this requires that **krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM** exist), and then use **those** credentials to obtain credentials for use in the **C.EXAMPLE.COM** realm (using **krbtgt/C.EXAMPLE.COM@B.EXAMPLE.COM**).

If those clients wish to contact a service in the **D.EXAMPLE.COM** realm, they will first need to obtain necessary credentials from the **B.EXAMPLE.COM** realm, and then credentials from the **C.EXAMPLE.COM** realm, before finally obtaining credentials for use with the **D.EXAMPLE.COM** realm.

Note

Without a **capath** entry indicating otherwise, Kerberos assumes that cross-realm trust relationships form a hierarchy.

Clients in the **A.EXAMPLE.COM** realm can obtain cross-realm credentials from **B.EXAMPLE.COM** realm directly. Without the "." indicating this, the client would instead attempt to use a hierarchical path, in this case:

```
A.EXAMPLE.COM → EXAMPLE.COM → B.EXAMPLE.COM
```

2.7.10. Additional Resources

For more information about Kerberos, refer to the following resources.

2.7.10.1. Installed Kerberos Documentation

- The *Kerberos V5 Installation Guide* and the *Kerberos V5 System Administrator's Guide* in PostScript and HTML formats. These can be found in the `/usr/share/doc/krb5-server-<version-number>/` directory (where `<version-number>` is the version number of the `krb5-server` package installed on your system).
- The *Kerberos V5 UNIX User's Guide* in PostScript and HTML formats. These can be found in the `/usr/share/doc/krb5-workstation-<version-number>/` directory (where `<version-number>` is the version number of the `krb5-workstation` package installed on your system).
- Kerberos man pages — There are a number of man pages for the various applications and configuration files involved with a Kerberos implementation. The following is a list of some of the more important man pages.

Client Applications

- **man `krberos`** — An introduction to the Kerberos system which describes how credentials work and provides recommendations for obtaining and destroying Kerberos tickets. The bottom of the man page references a number of related man pages.
- **man `kinit`** — Describes how to use this command to obtain and cache a ticket-granting ticket.
- **man `kdestroy`** — Describes how to use this command to destroy Kerberos credentials.
- **man `klist`** — Describes how to use this command to list cached Kerberos credentials.

Administrative Applications

- **man `kadmin`** — Describes how to use this command to administer the Kerberos V5 database.
- **man `kdb5_util`** — Describes how to use this command to create and perform low-level administrative functions on the Kerberos V5 database.

Server Applications

- **man `krb5kdc`** — Describes available command line options for the Kerberos V5 KDC.
- **man `kadmind`** — Describes available command line options for the Kerberos V5 administration server.

Configuration Files

- **man `krb5.conf`** — Describes the format and options available within the configuration file for the Kerberos V5 library.
- **man `kdc.conf`** — Describes the format and options available within the configuration file for the Kerberos V5 AS and KDC.

2.7.10.2. Useful Kerberos Websites

- <http://web.mit.edu/kerberos/www/> — *Kerberos: The Network Authentication Protocol* webpage from MIT.
- <http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html> — The Kerberos Frequently Asked Questions (FAQ).

- <ftp://athena-dist.mit.edu/pub/kerberos/doc/usenix.PS> — The PostScript version of *Kerberos: An Authentication Service for Open Network Systems* by Jennifer G. Steiner, Clifford Neuman, and Jeffrey I. Schiller. This document is the original paper describing Kerberos.
- <http://web.mit.edu/kerberos/www/dialogue.html> — *Designing an Authentication System: a Dialogue in Four Scenes* originally by Bill Bryant in 1988, modified by Theodore Ts'o in 1997. This document is a conversation between two developers who are thinking through the creation of a Kerberos-style authentication system. The conversational style of the discussion make this a good starting place for people who are completely unfamiliar with Kerberos.
- <http://www.ornl.gov/~jar/HowToKerb.html> — *How to Kerberize your site* is a good reference for kerberizing a network.
- <http://www.networkcomputing.com/netdesign/kerb1.html> — *Kerberos Network Design Manual* is a thorough overview of the Kerberos system.

2.8. Virtual Private Networks (VPNs)

Organizations with several satellite offices often connect to each other with dedicated lines for efficiency and protection of sensitive data in transit. For example, many businesses use frame relay or *Asynchronous Transfer Mode* (ATM) lines as an end-to-end networking solution to link one office with others. This can be an expensive proposition, especially for small to medium sized businesses (SMBs) that want to expand without paying the high costs associated with enterprise-level, dedicated digital circuits.

To address this need, *Virtual Private Networks* (VPNs) were developed. Following the same functional principles as dedicated circuits, VPNs allow for secured digital communication between two parties (or networks), creating a *Wide Area Network* (WAN) from existing *Local Area Networks* (LANs). Where it differs from frame relay or ATM is in its transport medium. VPNs transmit over IP using datagrams as the transport layer, making it a secure conduit through the Internet to an intended destination. Most free software VPN implementations incorporate open standard encryption methods to further mask data in transit.

Some organizations employ hardware VPN solutions to augment security, while others use software or protocol-based implementations. Several vendors provide hardware VPN solutions, such as Cisco, Nortel, IBM, and Checkpoint. There is a free software-based VPN solution for Linux called FreeS/Wan that utilizes a standardized *Internet Protocol Security* (IPsec) implementation. These VPN solutions, irrespective of whether they are hardware or software based, act as specialized routers that exist between the IP connection from one office to another.

2.8.1. How Does a VPN Work?

When a packet is transmitted from a client, it sends it through the VPN router or gateway, which adds an *Authentication Header* (AH) for routing and authentication. The data is then encrypted and, finally, enclosed with an *Encapsulating Security Payload* (ESP). This latter constitutes the decryption and handling instructions.

The receiving VPN router strips the header information, decrypts the data, and routes it to its intended destination (either a workstation or other node on a network). Using a network-to-network connection, the receiving node on the local network receives the packets already decrypted and ready for processing. The encryption/decryption process in a network-to-network VPN connection is transparent to a local node.

With such a heightened level of security, an attacker must not only intercept a packet, but decrypt the packet as well. Intruders who employ a man-in-the-middle attack between a server and client must also have access to at least one of the private keys for authenticating sessions. Because they employ several layers of authentication and encryption, VPNs are a secure and effective means of connecting multiple remote nodes to act as a unified intranet.

2.8.2. VPNs and Fedora

Fedora provides various options in terms of implementing a software solution to securely connect to a WAN. *Internet Protocol Security* (IPsec) is the supported VPN implementation for Fedora, and sufficiently addresses the usability needs of organizations with branch offices or remote users.

2.8.3. IPsec

Fedora supports IPsec for connecting remote hosts and networks to each other using a secure tunnel on a common carrier network such as the Internet. IPsec can be implemented using a host-to-host (one computer workstation to another) or network-to-network (one LAN/WAN to another) configuration.

The IPsec implementation in Fedora uses *Internet Key Exchange* (IKE), a protocol implemented by the Internet Engineering Task Force (IETF), used for mutual authentication and secure associations between connecting systems.

2.8.4. Creating an IPsec Connection

An IPsec connection is split into two logical phases. In phase 1, an IPsec node initializes the connection with the remote node or network. The remote node or network checks the requesting node's credentials and both parties negotiate the authentication method for the connection.

On Fedora systems, an IPsec connection uses the *pre-shared key* method of IPsec node authentication. In a pre-shared key IPsec connection, both hosts must use the same key in order to move to Phase 2 of the IPsec connection.

Phase 2 of the IPsec connection is where the *Security Association* (SA) is created between IPsec nodes. This phase establishes an SA database with configuration information, such as the encryption method, secret session key exchange parameters, and more. This phase manages the actual IPsec connection between remote nodes and networks.

The Fedora implementation of IPsec uses IKE for sharing keys between hosts across the Internet. The **racoon** keying daemon handles the IKE key distribution and exchange. Refer to the **racoon** man page for more information about this daemon.

2.8.5. IPsec Installation

Implementing IPsec requires that the **ipsec-tools** RPM package be installed on all IPsec hosts (if using a host-to-host configuration) or routers (if using a network-to-network configuration). The RPM package contains essential libraries, daemons, and configuration files for setting up the IPsec connection, including:

- **/sbin/setkey** — manipulates the key management and security attributes of IPsec in the kernel. This executable is controlled by the **racoon** key management daemon. Refer to the **setkey(8)** man page for more information.
- **/usr/sbin/racoon** — the IKE key management daemon, used to manage and control security associations and key sharing between IPsec-connected systems.

- `/etc/racoon/racoon.conf` — the **racoon** daemon configuration file used to configure various aspects of the IPsec connection, including authentication methods and encryption algorithms used in the connection. Refer to the `racoon.conf(5)` man page for a complete listing of available directives.

To configure IPsec on Fedora, you can use the **Network Administration Tool**, or manually edit the networking and IPsec configuration files.

- To connect two network-connected hosts via IPsec, refer to [Section 2.8.6, “IPsec Host-to-Host Configuration”](#).
- To connect one LAN/WAN to another via IPsec, refer to [Section 2.8.7, “IPsec Network-to-Network Configuration”](#).

2.8.6. IPsec Host-to-Host Configuration

IPsec can be configured to connect one desktop or workstation (host) to another using a host-to-host connection. This type of connection uses the network to which each host is connected to create a secure tunnel between each host. The requirements of a host-to-host connection are minimal, as is the configuration of IPsec on each host. The hosts need only a dedicated connection to a carrier network (such as the Internet) and Fedora to create the IPsec connection.

2.8.6.1. Host-to-Host Connection

A host-to-host IPsec connection is an encrypted connection between two systems, both running IPsec with the same authentication key. With the IPsec connection active, any network traffic between the two hosts is encrypted.

To configure a host-to-host IPsec connection, use the following steps for each host:



Note

You should perform the following procedures on the actual machine that you are configuring. Avoid attempting to configure and establish IPsec connections remotely.

1. In a command shell, type **system-config-network** to start the **Network Administration Tool**.
2. On the **IPsec** tab, click **New** to start the IPsec configuration wizard.
3. Click **Forward** to start configuring a host-to-host IPsec connection.
4. Enter a unique name for the connection, for example, **ipsec0**. If required, select the check box to automatically activate the connection when the computer starts. Click **Forward** to continue.
5. Select **Host to Host encryption** as the connection type, and then click **Forward**.
6. Select the type of encryption to use: manual or automatic.

If you select manual encryption, an encryption key must be provided later in the process. If you select automatic encryption, the **racoon** daemon manages the encryption key. The **ipsec-tools** package must be installed if you want to use automatic encryption.

Click **Forward** to continue.

7. Enter the IP address of the remote host.

To determine the IP address of the remote host, use the following command *on the remote host*:


```
[root@myServer ~] # /sbin/ifconfig <device>
```

where *<device>* is the Ethernet device that you want to use for the VPN connection.

If only one Ethernet card exists in the system, the device name is typically `eth0`. The following example shows the relevant information from this command (note that this is an example output only):

```
eth0      Link encap:Ethernet  HWaddr 00:0C:6E:E8:98:1D  
          inet addr:172.16.44.192  Bcast:172.16.45.255  Mask:255.255.254.0
```

The IP address is the number following the `inet addr:` label.



Note

For host-to-host connections, both hosts should have a public, routable address. Alternatively, both hosts can have a private, non-routable address (for example, from the 10.x.x.x or 192.168.x.x ranges) as long as they are on the same LAN.

If the hosts are on different LANs, or one has a public address while the other has a private address, refer to *Section 2.8.7, “IPsec Network-to-Network Configuration”*.

Click **Forward** to continue.

8. If manual encryption was selected in step 6, specify the encryption key to use, or click **Generate** to create one.
 - a. Specify an authentication key or click **Generate** to generate one. It can be any combination of numbers and letters.
 - b. Click **Forward** to continue.
9. Verify the information on the **IPsec — Summary** page, and then click **Apply**.
10. Click **File => Save** to save the configuration.

You may need to restart the network for the changes to take effect. To restart the network, use the following command:

```
[root@myServer ~]# service network restart
```

11. Select the IPsec connection from the list and click the **Activate** button.
12. Repeat the entire procedure for the other host. It is essential that the same keys from step 8 be used on the other hosts. Otherwise, IPsec will not work.

After configuring the IPsec connection, it appears in the IPsec list as shown in [Figure 2.10, “IPsec Connection”](#).

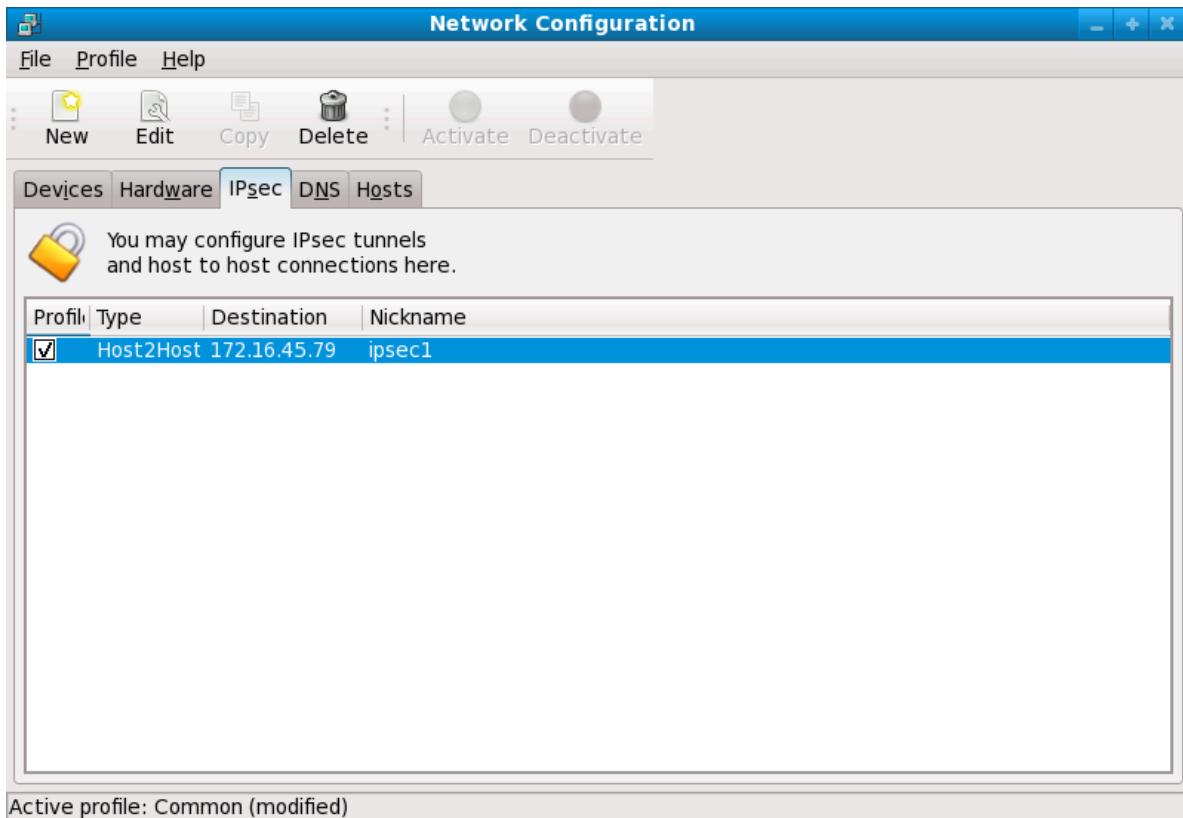


Figure 2.10. IPsec Connection

The following files are created when the IPsec connection is configured:

- `/etc/sysconfig/network-scripts/ifcfg-<nickname>`
- `/etc/sysconfig/network-scripts/keys-<nickname>`
- `/etc/racoon/<remote-ip>.conf`
- `/etc/racoon/psk.txt`

If automatic encryption is selected, `/etc/racoon/racoon.conf` is also created.

When the interface is up, `/etc/racoon/racoon.conf` is modified to include `<remote-ip>.conf`.

2.8.6.2. Manual IPsec Host-to-Host Configuration

The first step in creating a connection is to gather system and network information from each workstation. For a host-to-host connection, you need the following:

- The IP address of each host
- A unique name, for example, `ipsec1`. This is used to identify the IPsec connection and to distinguish it from other devices or connections.
- A fixed encryption key or one automatically generated by **racoon**.
- A pre-shared authentication key that is used during the initial stage of the connection and to exchange encryption keys during the session.

For example, suppose Workstation A and Workstation B want to connect to each other through an IPsec tunnel. They want to connect using a pre-shared key with the value of `Key_Value01`, and the users agree to let **racoon** automatically generate and share an authentication key between each host. Both host users decide to name their connections `ipsec1`.



Note

You should choose a PSK that uses a mixture of upper- and lower-case characters, numbers and punctuation. An easily-guessable PSK constitutes a security risk.

It is not necessary to use the same connection name for each host. You should choose a name that is convenient and meaningful for your installation.

The following is the IPsec configuration file for Workstation A for a host-to-host IPsec connection with Workstation B. The unique name to identify the connection in this example is `ipsec1`, so the resulting file is called `/etc/sysconfig/network-scripts/ifcfg-ipsec1`.

```
DST=X.X.X.XTYPE=IPSEC
ONBOOT=no
IKE_METHOD=PSK
```

For Workstation A, `X.X.X.X` is the IP address of Workstation B. For Workstation B, `X.X.X.X` is the IP address of Workstation A. This connection is not set to initiate on boot-up (`ONBOOT=no`) and it uses the pre-shared key method of authentication (`IKE_METHOD=PSK`).

The following is the content of the pre-shared key file (called `/etc/sysconfig/network-scripts/keys-ipsec1`) that both workstations need to authenticate each other. The contents of this file should be identical on both workstations, and only the root user should be able to read or write this file.

```
IKE_PSK=Key_Value01
```



Important

To change the `keys-ipsec1` file so that only the root user can read or edit the file, use the following command after creating the file:

```
[root@myServer ~] # chmod 600 /etc/sysconfig/network-scripts/keys-ipsec1
```

To change the authentication key at any time, edit the `keys-ipsec1` file on both workstations. *Both authentication keys must be identical for proper connectivity.*

The next example shows the specific configuration for the phase 1 connection to the remote host. The file is called `X.X.X.X.conf`, where `X.X.X.X` is the IP address of the remote IPsec host. Note that this file is automatically generated when the IPsec tunnel is activated and should not be edited directly.

```
remote X.X.X.X{
    exchange_mode aggressive, main;
    my_identifier address;
    proposal {
```

```
encryption_algorithm 3des;
hash_algorithm sha1;
authentication_method pre_shared_key;
dh_group 2 ;
}
}
```

The default phase 1 configuration file that is created when an IPsec connection is initialized contains the following statements used by the Fedora implementation of IPsec:

`remote X.X.X.X`

Specifies that the subsequent stanzas of this configuration file apply only to the remote node identified by the `X.X.X.X` IP address.

`exchange_mode aggressive`

The default configuration for IPsec on Fedora uses an aggressive authentication mode, which lowers the connection overhead while allowing configuration of several IPsec connections with multiple hosts.

`my_identifier address`

Specifies the identification method to use when authenticating nodes. Fedora uses IP addresses to identify nodes.

`encryption_algorithm 3des`

Specifies the encryption cipher used during authentication. By default, *Triple Data Encryption Standard* (3DES) is used.

`hash_algorithm sha1;`

Specifies the hash algorithm used during phase 1 negotiation between nodes. By default, Secure Hash Algorithm version 1 is used.

`authentication_method pre_shared_key`

Specifies the authentication method used during node negotiation. By default, Fedora uses pre-shared keys for authentication.

`dh_group 2`

Specifies the Diffie-Hellman group number for establishing dynamically-generated session keys. By default, modp1024 (group 2) is used.

2.8.6.2.1. The Racoon Configuration File

The `/etc/racoon/racoon.conf` files should be identical on all IPsec nodes *except* for the **include** `"/etc/racoon/X.X.X.X.conf"` statement. This statement (and the file it references) is generated when the IPsec tunnel is activated. For Workstation A, the `X.X.X.X` in the **include** statement is Workstation B's IP address. The opposite is true of Workstation B. The following shows a typical **racoon.conf** file when the IPsec connection is activated.

```
# Racoon IKE daemon configuration file.
# See 'man racoon.conf' for a description of the format and entries.

path include "/etc/racoon";
path pre_shared_key "/etc/racoon/psk.txt";
path certificate "/etc/racoon/certs";

sainfo anonymous
{
```

```

    pfs_group 2;
    lifetime time 1 hour ;
    encryption_algorithm 3des, blowfish 448, rijndael ;
    authentication_algorithm hmac_sha1, hmac_md5 ;
    compression_algorithm deflate ;
}
include "/etc/racoon/X.X.X.X.conf";

```

This default **racoon.conf** file includes defined paths for IPsec configuration, pre-shared key files, and certificates. The fields in `sainfo anonymous` describe the phase 2 SA between the IPsec nodes — the nature of the IPsec connection (including the supported encryption algorithms used) and the method of exchanging keys. The following list defines the fields of phase 2:

sainfo anonymous

Denotes that SA can anonymously initialize with any peer provided that the IPsec credentials match.

pfs_group 2

Defines the Diffie-Hellman key exchange protocol, which determines the method by which the IPsec nodes establish a mutual temporary session key for the second phase of IPsec connectivity. By default, the Fedora implementation of IPsec uses group 2 (or `modp1024`) of the Diffie-Hellman cryptographic key exchange groups. Group 2 uses a 1024-bit modular exponentiation that prevents attackers from decrypting previous IPsec transmissions even if a private key is compromised.

lifetime time 1 hour

This parameter specifies the lifetime of an SA and can be quantified either by time or by bytes of data. The default Fedora implementation of IPsec specifies a one hour lifetime.

encryption_algorithm 3des, blowfish 448, rijndael

Specifies the supported encryption ciphers for phase 2. Fedora supports 3DES, 448-bit Blowfish, and Rijndael (the cipher used in the *Advanced Encryption Standard*, or AES).

authentication_algorithm hmac_sha1, hmac_md5

Lists the supported hash algorithms for authentication. Supported modes are sha1 and md5 hashed message authentication codes (HMAC).

compression_algorithm deflate

Defines the Deflate compression algorithm for IP Payload Compression (IPCOMP) support, which allows for potentially faster transmission of IP datagrams over slow connections.

To start the connection, use the following command on each host:

```
[root@myServer ~]# /sbin/ifup <nickname>
```

where `<nickname>` is the name you specified for the IPsec connection.

To test the IPsec connection, run the **tcpdump** utility to view the network packets being transferred between the hosts and verify that they are encrypted via IPsec. The packet should include an AH header and should be shown as ESP packets. ESP means it is encrypted. For example:

```
[root@myServer ~]# tcpdump -n -i eth0 host <targetSystem>
IP 172.16.45.107 > 172.16.44.192: AH(spi=0x0954ccb6,seq=0xbb): ESP(spi=0x0c9f2164,seq=0xbb)
```

2.8.7. IPsec Network-to-Network Configuration

IPsec can also be configured to connect an entire network (such as a LAN or WAN) to a remote network using a network-to-network connection. A network-to-network connection requires the setup of IPsec routers on each side of the connecting networks to transparently process and route information from one node on a LAN to a node on a remote LAN. *Figure 2.11, “A network-to-network IPsec tunneled connection”* shows a network-to-network IPsec tunneled connection.




Figure 2.11. A network-to-network IPsec tunneled connection

This diagram shows two separate LANs separated by the Internet. These LANs use IPsec routers to authenticate and initiate a connection using a secure tunnel through the Internet. Packets that are intercepted in transit would require brute-force decryption in order to crack the cipher protecting the packets between these LANs. The process of communicating from one node in the 192.168.1.0/24 IP range to another in the 192.168.2.0/24 range is completely transparent to the nodes as the processing, encryption/decryption, and routing of the IPsec packets are completely handled by the IPsec router.

The information needed for a network-to-network connection include:

- The externally-accessible IP addresses of the dedicated IPsec routers
- The network address ranges of the LAN/WAN served by the IPsec routers (such as 192.168.1.0/24 or 10.0.1.0/24)
- The IP addresses of the gateway devices that route the data from the network nodes to the Internet
- A unique name, for example, `ipsec1`. This is used to identify the IPsec connection and to distinguish it from other devices or connections.
- A fixed encryption key or one automatically generated by **racoon**
- A pre-shared authentication key that is used during the initial stage of the connection and to exchange encryption keys during the session.

2.8.7.1. Network-to-Network (VPN) Connection

A network-to-network IPsec connection uses two IPsec routers, one for each network, through which the network traffic for the private subnets is routed.

For example, as shown in *Figure 2.12, “Network-to-Network IPsec”*, if the 192.168.1.0/24 private network sends network traffic to the 192.168.2.0/24 private network, the packets go through gateway0, to ipsec0, through the Internet, to ipsec1, to gateway1, and to the 192.168.2.0/24 subnet.

IPsec routers require publicly addressable IP addresses and a second Ethernet device connected to their respective private networks. Traffic only travels through an IPsec router if it is intended for another IPsec router with which it has an encrypted connection.




Figure 2.12. Network-to-Network IPsec

Alternate network configuration options include a firewall between each IP router and the Internet, and an intranet firewall between each IPsec router and subnet gateway. The IPsec router and the gateway for the subnet can be one system with two Ethernet devices: one with a public IP address that acts as the IPsec router; and one with a private IP address that acts as the gateway for the private subnet. Each IPsec router can use the gateway for its private network or a public gateway to send the packets to the other IPsec router.

Use the following procedure to configure a network-to-network IPsec connection:

1. In a command shell, type **system-config-network** to start the **Network Administration Tool**.
2. On the **IPsec** tab, click **New** to start the IPsec configuration wizard.
3. Click **Forward** to start configuring a network-to-network IPsec connection.
4. Enter a unique nickname for the connection, for example, **ipsec0**. If required, select the check box to automatically activate the connection when the computer starts. Click **Forward** to continue.
5. Select **Network to Network encryption (VPN)** as the connection type, and then click **Forward**.
6. Select the type of encryption to use: manual or automatic.

If you select manual encryption, an encryption key must be provided later in the process. If you select automatic encryption, the **racoon** daemon manages the encryption key. The **ipsec-tools** package must be installed if you want to use automatic encryption.

Click **Forward** to continue.

7. On the **Local Network** page, enter the following information:
 - **Local Network Address** — The IP address of the device on the IPsec router connected to the private network.
 - **Local Subnet Mask** — The subnet mask of the local network IP address.
 - **Local Network Gateway** — The gateway for the private subnet.

Click **Forward** to continue.

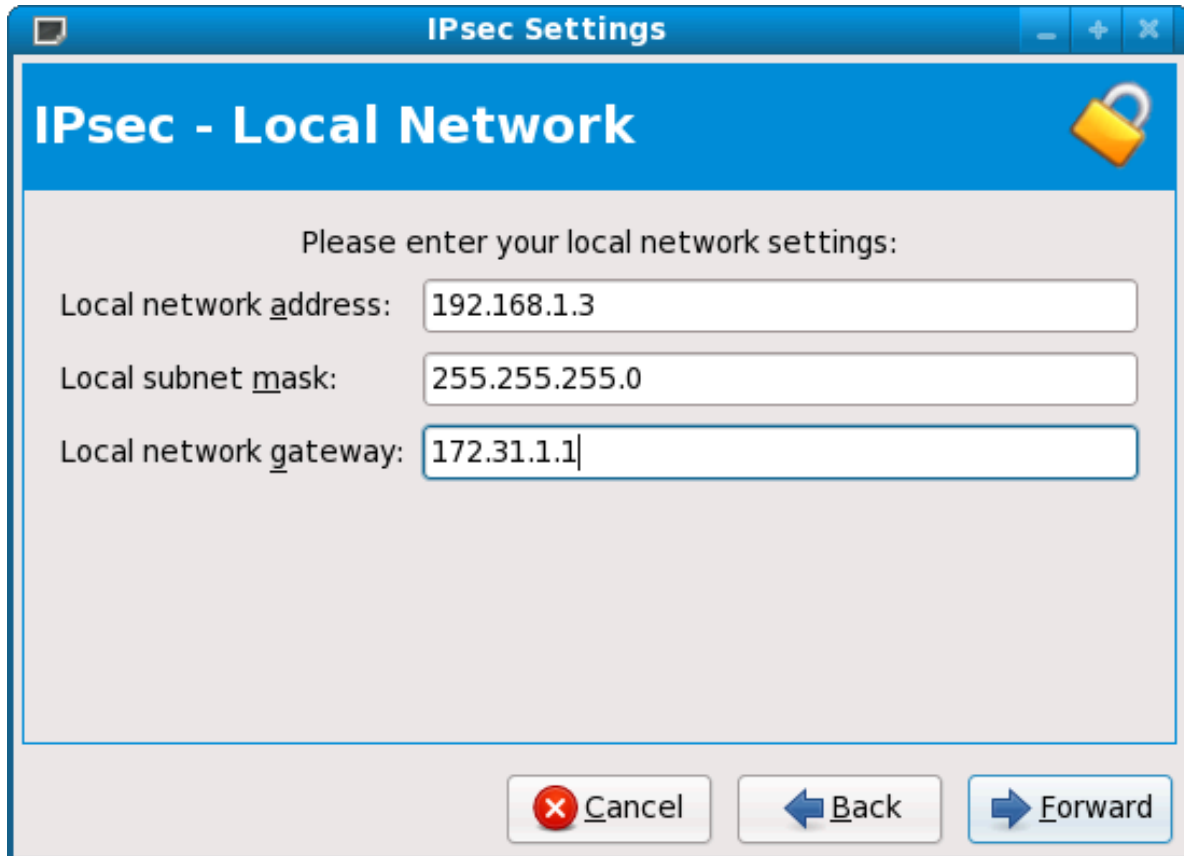


Figure 2.13. Local Network Information

8. On the **Remote Network** page, enter the following information:
 - **Remote IP Address** — The publicly addressable IP address of the IPsec router for the *other* private network. In our example, for ipsec0, enter the publicly addressable IP address of ipsec1, and vice versa.
 - **Remote Network Address** — The network address of the private subnet behind the *other* IPsec router. In our example, enter **192 . 168 . 1 . 0** if configuring ipsec1, and enter **192 . 168 . 2 . 0** if configuring ipsec0.
 - **Remote Subnet Mask** — The subnet mask of the remote IP address.
 - **Remote Network Gateway** — The IP address of the gateway for the remote network address.
 - If manual encryption was selected in step 6, specify the encryption key to use or click **Generate** to create one.

Specify an authentication key or click **Generate** to generate one. This key can be any combination of numbers and letters.

Click **Forward** to continue.

The screenshot shows a window titled "IPsec Settings" with a sub-header "IPsec - Remote Network". Below the sub-header is a lock icon. The main content area contains the text "Please enter your remote network settings:" followed by four input fields:

- Remote IP address: 172.16.57.27
- Remote network address: 192.168.1.0
- Remote subnet mask: 255.255.255.0
- Remote network gateway: 192.168.1.1

At the bottom of the window are three buttons: "Cancel" (with a red X icon), "Back" (with a left arrow icon), and "Forward" (with a right arrow icon).

Figure 2.14. Remote Network Information

9. Verify the information on the **IPsec — Summary** page, and then click **Apply**.
10. Select **File => Save** to save the configuration.
11. Select the IPsec connection from the list, and then click **Activate** to activate the connection.
12. Enable IP forwarding:
 - a. Edit `/etc/sysctl.conf` and set `net.ipv4.ip_forward` to **1**.
 - b. Use the following command to enable the change:

```
[root@myServer ~]# /sbin/sysctl -p /etc/sysctl.conf
```

The network script to activate the IPsec connection automatically creates network routes to send packets through the IPsec router if necessary.

2.8.7.2. Manual IPsec Network-to-Network Configuration

Suppose LAN A (lana.example.com) and LAN B (lanb.example.com) want to connect to each other through an IPsec tunnel. The network address for LAN A is in the 192.168.1.0/24 range, while LAN B uses the 192.168.2.0/24 range. The gateway IP address is 192.168.1.254 for LAN A and 192.168.2.254 for LAN B. The IPsec routers are separate from each LAN gateway and use two network devices: eth0 is assigned to an externally-accessible static IP address which accesses the

Internet, while eth1 acts as a routing point to process and transmit LAN packets from one network node to the remote network nodes.

The IPsec connection between each network uses a pre-shared key with the value of `r3dh4t11nux`, and the administrators of A and B agree to let **racoon** automatically generate and share an authentication key between each IPsec router. The administrator of LAN A decides to name the IPsec connection `ipsec0`, while the administrator of LAN B names the IPsec connection `ipsec1`.

The following example shows the contents of the `ifcfg` file for a network-to-network IPsec connection for LAN A. The unique name to identify the connection in this example is `ipsec0`, so the resulting file is called `/etc/sysconfig/network-scripts/ifcfg-ipsec0`.

```
TYPE=IPSEC
ONBOOT=yes
IKE_METHOD=PSK
SRCGW=192.168.1.254
DSTGW=192.168.2.254
SRCNET=192.168.1.0/24
DSTNET=192.168.2.0/24
DST=X.X.X.X
```

The following list describes the contents of this file:

`TYPE=IPSEC`

Specifies the type of connection.

`ONBOOT=yes`

Specifies that the connection should initiate on boot-up.

`IKE_METHOD=PSK`

Specifies that the connection uses the pre-shared key method of authentication.

`SRCGW=192.168.1.254`

The IP address of the source gateway. For LAN A, this is the LAN A gateway, and for LAN B, the LAN B gateway.

`DSTGW=192.168.2.254`

The IP address of the destination gateway. For LAN A, this is the LAN B gateway, and for LAN B, the LAN A gateway.

`SRCNET=192.168.1.0/24`

Specifies the source network for the IPsec connection, which in this example is the network range for LAN A.

`DSTNET=192.168.2.0/24`

Specifies the destination network for the IPsec connection, which in this example is the network range for LAN B.

`DST=X.X.X.X`

The externally-accessible IP address of LAN B.

The following example is the content of the pre-shared key file called `/etc/sysconfig/network-scripts/keys-ipsecX` (where X is 0 for LAN A and 1 for LAN B) that both networks use to authenticate each other. The contents of this file should be identical and only the root user should be able to read or write this file.


```
IKE_PSK=r3dh4t11nux
```



Important

To change the **keys-ipsecX** file so that only the root user can read or edit the file, use the following command after creating the file:

```
chmod 600 /etc/sysconfig/network-scripts/keys-ipsec1
```

To change the authentication key at any time, edit the **keys-ipsecX** file on both IPsec routers. *Both keys must be identical for proper connectivity.*

The following example is the contents of the **/etc/racoon/racoon.conf** configuration file for the IPsec connection. Note that the `include` line at the bottom of the file is automatically generated and only appears if the IPsec tunnel is running.

```
# Racoon IKE daemon configuration file.
# See 'man racoon.conf' for a description of the format and entries.
path include "/etc/racoon";
path pre_shared_key "/etc/racoon/psk.txt";
path certificate "/etc/racoon/certs";

sainfo anonymous
{
  pfs_group 2;
  lifetime time 1 hour ;
  encryption_algorithm 3des, blowfish 448, rijndael ;
  authentication_algorithm hmac_sha1, hmac_md5 ;
  compression_algorithm deflate ;
}
include "/etc/racoon/X.X.X.X.conf"
```

The following is the specific configuration for the connection to the remote network. The file is called **X.X.X.X.conf** (where X.X.X.X is the IP address of the remote IPsec router). Note that this file is automatically generated when the IPsec tunnel is activated and should not be edited directly.

```
remote X.X.X.X{
  exchange_mode aggressive, main;
  my_identifier address;
  proposal {
    encryption_algorithm 3des;
    hash_algorithm sha1;
    authentication_method pre_shared_key;
    dh_group 2 ;
  }
}
```

Prior to starting the IPsec connection, IP forwarding should be enabled in the kernel. To enable IP forwarding:

1. Edit **/etc/sysctl.conf** and set `net.ipv4.ip_forward` to **1**.
2. Use the following command to enable the change:

```
[root@myServer ~] # sysctl -p /etc/sysctl.conf
```

To start the IPsec connection, use the following command on each router:

```
[root@myServer ~] # /sbin/ufup ipsec0
```

The connections are activated, and both LAN A and LAN B are able to communicate with each other. The routes are created automatically via the initialization script called by running **ufup** on the IPsec connection. To show a list of routes for the network, use the following command:

```
[root@myServer ~] # /sbin/ip route list
```

To test the IPsec connection, run the **tcpdump** utility on the externally-routable device (eth0 in this example) to view the network packets being transferred between the hosts (or networks), and verify that they are encrypted via IPsec. For example, to check the IPsec connectivity of LAN A, use the following command:

```
[root@myServer ~] # tcpdump -n -i eth0 host lana.example.com
```

The packet should include an AH header and should be shown as ESP packets. ESP means it is encrypted. For example (back slashes denote a continuation of one line):

```
12:24:26.155529 lanb.example.com > lana.example.com: AH(spi=0x021c9834,seq=0x358): \
lanb.example.com > lana.example.com: ESP(spi=0x00c887ad,seq=0x358) (DF) \
(ipip-proto-4)
```

2.8.8. Starting and Stopping an IPsec Connection

If the IPsec connection was not configured to activate on boot, you can control it from the command line.

To start the connection, use the following command on each host for host-to-host IPsec, or each IPsec router for network-to-network IPsec:

```
[root@myServer ~] # /sbin/ufup <nickname>
```

where *<nickname>* is the nickname configured earlier, such as ipsec0.

To stop the connection, use the following command:

```
[root@myServer ~] # /sbin/ufdown <nickname>
```

2.9. Firewalls

Information security is commonly thought of as a process and not a product. However, standard security implementations usually employ some form of dedicated mechanism to control access privileges and restrict network resources to users who are authorized, identifiable, and traceable. Fedora includes several tools to assist administrators and security engineers with network-level access control issues.

Firewalls are one of the core components of a network security implementation. Several vendors market firewall solutions catering to all levels of the marketplace: from home users protecting one PC to data center solutions safeguarding vital enterprise information. Firewalls can be stand-alone hardware solutions, such as firewall appliances by Cisco, Nokia, and Sonicwall. Vendors such as

Checkpoint, McAfee, and Symantec have also developed proprietary software firewall solutions for home and business markets.

Apart from the differences between hardware and software firewalls, there are also differences in the way firewalls function that separate one solution from another. [Table 2.2, “Firewall Types”](#) details three common types of firewalls and how they function:

Method	Description	Advantages	Disadvantages
NAT	<i>Network Address Translation</i> (NAT) places private IP subnetworks behind one or a small pool of public IP addresses, masquerading all requests to one source rather than several. The Linux kernel has built-in NAT functionality through the Netfilter kernel subsystem.	<ul style="list-style-type: none"> · Can be configured transparently to machines on a LAN · Protection of many machines and services behind one or more external IP addresses simplifies administration duties · Restriction of user access to and from the LAN can be configured by opening and closing ports on the NAT firewall/gateway 	<ul style="list-style-type: none"> · Cannot prevent malicious activity once users connect to a service outside of the firewall
Packet Filter	A packet filtering firewall reads each data packet that passes through a LAN. It can read and process packets by header information and filters the packet based on sets of programmable rules implemented by the firewall administrator. The Linux kernel has built-in packet filtering functionality through the Netfilter kernel subsystem.	<ul style="list-style-type: none"> · Customizable through the iptables front-end utility · Does not require any customization on the client side, as all network activity is filtered at the router level rather than the application level · Since packets are not transmitted through a proxy, network performance is faster due to direct connection from client to remote host 	<ul style="list-style-type: none"> · Cannot filter packets for content like proxy firewalls · Processes packets at the protocol layer, but cannot filter packets at an application layer · Complex network architectures can make establishing packet filtering rules difficult, especially if coupled with <i>IP masquerading</i> or local subnets and DMZ networks
Proxy	Proxy firewalls filter all requests of a certain protocol or type from LAN clients to a proxy machine, which then makes those requests to the Internet on behalf of the local client. A proxy machine acts as a buffer between malicious remote users and the internal network client machines.	<ul style="list-style-type: none"> · Gives administrators control over what applications and protocols function outside of the LAN · Some proxy servers can cache frequently-accessed data locally rather than having to use the Internet connection to request it. This helps to reduce bandwidth consumption · Proxy services can be logged and monitored closely, allowing tighter control over resource utilization on the network 	<ul style="list-style-type: none"> · Proxies are often application-specific (HTTP, Telnet, etc.), or protocol-restricted (most proxies work with TCP-connected services only) · Application services cannot run behind a proxy, so your application servers must use a separate form of network security · Proxies can become a network bottleneck, as all requests and transmissions are passed through one source rather than directly

Method	Description	Advantages	Disadvantages
			from a client to a remote service

Table 2.2. Firewall Types

2.9.1. Netfilter and IPTables

The Linux kernel features a powerful networking subsystem called *Netfilter*. The Netfilter subsystem provides stateful or stateless packet filtering as well as NAT and IP masquerading services. Netfilter also has the ability to *mangle* IP header information for advanced routing and connection state management. Netfilter is controlled using the **iptables** tool.

2.9.1.1. IPTables Overview

The power and flexibility of Netfilter is implemented using the **iptables** administration tool, a command line tool similar in syntax to its predecessor, **ipchains**, which Netfilter/iptables replaced in the Linux kernel 2.4 and above.

iptables uses the Netfilter subsystem to enhance network connection, inspection, and processing. **iptables** features advanced logging, pre- and post-routing actions, network address translation, and port forwarding, all in one command line interface.

This section provides an overview of **iptables**. For more detailed information, refer to [Section 2.10, “IPTables”](#).

2.9.2. Basic Firewall Configuration

Just as a firewall in a building attempts to prevent a fire from spreading, a computer firewall attempts to prevent malicious software from spreading to your computer. It also helps to prevent unauthorized users from accessing your computer.

In a default Fedora installation, a firewall exists between your computer or network and any untrusted networks, for example the Internet. It determines which services on your computer remote users can access. A properly configured firewall can greatly increase the security of your system. It is recommended that you configure a firewall for any Fedora system with an Internet connection.

2.9.2.1. Firewall Configuration Tool

During the **Firewall Configuration** screen of the Fedora installation, you were given the option to enable a basic firewall as well as to allow specific devices, incoming services, and ports.

After installation, you can change this preference by using the **Firewall Configuration Tool**.

To start this application, use the following command:

```
[root@myServer ~] # system-config-firewall
```

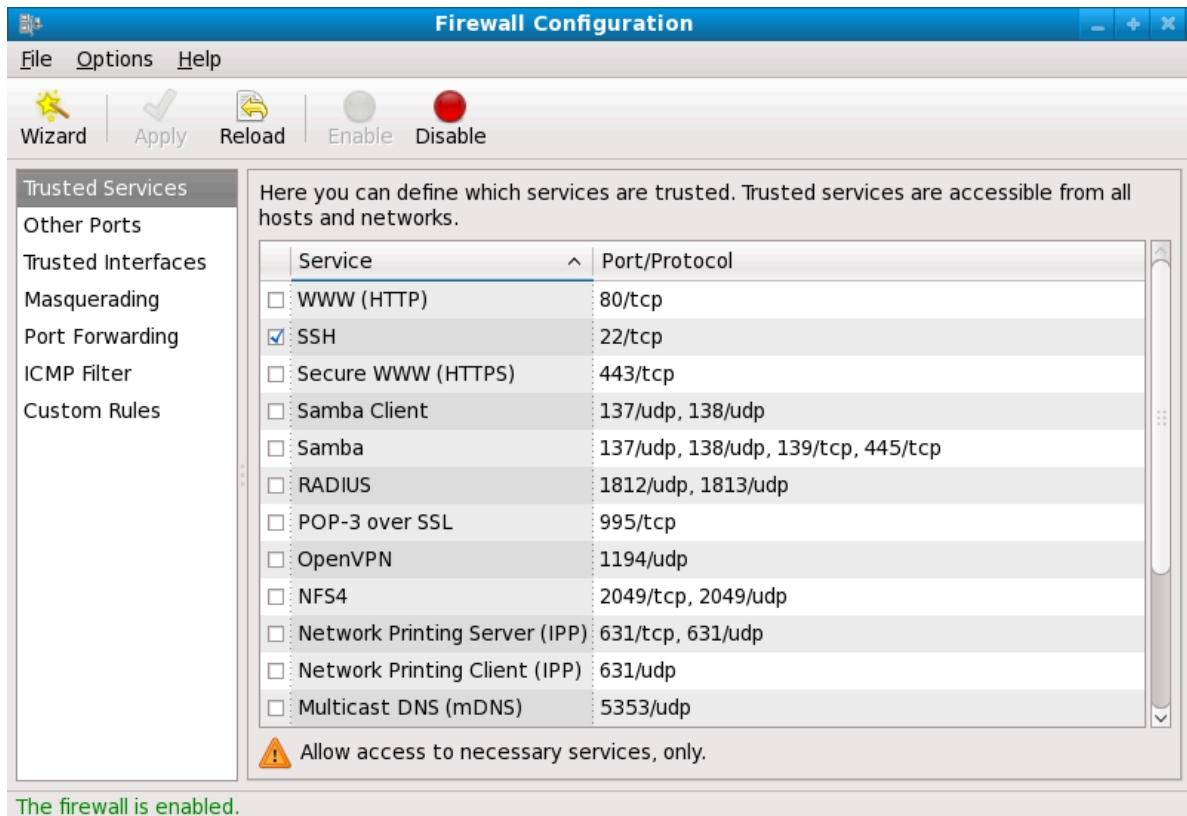


Figure 2.15. Firewall Configuration Tool

Note

The **Firewall Configuration Tool** only configures a basic firewall. If the system needs more complex rules, refer to *Section 2.10, “IPTables”* for details on configuring specific `iptables` rules.

2.9.2.2. Enabling and Disabling the Firewall

Select one of the following options for the firewall:

- **Disabled** — Disabling the firewall provides complete access to your system and does no security checking. This should only be selected if you are running on a trusted network (not the Internet) or need to configure a custom firewall using the `iptables` command line tool.



Warning

Firewall configurations and any customized firewall rules are stored in the `/etc/sysconfig/iptables` file. If you choose **Disabled** and click **OK**, these configurations and firewall rules will be lost.

- **Enabled** — This option configures the system to reject incoming connections that are not in response to outbound requests, such as DNS replies or DHCP requests. If access to services running on this machine is needed, you can choose to allow specific services through the firewall.

If you are connecting your system to the Internet, but do not plan to run a server, this is the safest choice.

2.9.2.3. Trusted Services

Enabling options in the **Trusted services** list allows the specified service to pass through the firewall.

WWW (HTTP)

The HTTP protocol is used by Apache (and by other Web servers) to serve web pages. If you plan on making your Web server publicly available, select this check box. This option is not required for viewing pages locally or for developing web pages. This service requires that the **httpd** package be installed.

Enabling **WWW (HTTP)** will not open a port for HTTPS, the SSL version of HTTP. If this service is required, select the **Secure WWW (HTTPS)** check box.

FTP

The FTP protocol is used to transfer files between machines on a network. If you plan on making your FTP server publicly available, select this check box. This service requires that the **vsftpd** package be installed.

SSH

Secure Shell (SSH) is a suite of tools for logging into and executing commands on a remote machine. To allow remote access to the machine via ssh, select this check box. This service requires that the **openssh-server** package be installed.

Telnet

Telnet is a protocol for logging into remote machines. Telnet communications are unencrypted and provide no security from network snooping. Allowing incoming Telnet access is not recommended. To allow remote access to the machine via telnet, select this check box. This service requires that the **telnet-server** package be installed.

Mail (SMTP)

SMTP is a protocol that allows remote hosts to connect directly to your machine to deliver mail. You do not need to enable this service if you collect your mail from your ISP's server using POP3 or IMAP, or if you use a tool such as **fetchmail**. To allow delivery of mail to your machine, select this check box. Note that an improperly configured SMTP server can allow remote machines to use your server to send spam.

NFS4

The Network File System (NFS) is a file sharing protocol commonly used on *NIX systems. Version 4 of this protocol is more secure than its predecessors. If you want to share files or directories on your system with other network users, select this check box.

Samba

Samba is an implementation of Microsoft's proprietary SMB networking protocol. If you need to share files, directories, or locally-connected printers with Microsoft Windows machines, select this check box.

2.9.2.4. Other Ports

The **Firewall Configuration Tool** includes an **Other ports** section for specifying custom IP ports as being trusted by **iptables**. For example, to allow IRC and Internet printing protocol (IPP) to pass through the firewall, add the following to the **Other ports** section:

```
194:tcp,631:tcp
```

2.9.2.5. Saving the Settings

Click **OK** to save the changes and enable or disable the firewall. If **Enable firewall** was selected, the options selected are translated to **iptables** commands and written to the `/etc/sysconfig/iptables` file. The **iptables** service is also started so that the firewall is activated immediately after saving the selected options. If **Disable firewall** was selected, the `/etc/sysconfig/iptables` file is removed and the **iptables** service is stopped immediately.

The selected options are also written to the `/etc/sysconfig/system-config-securitylevel` file so that the settings can be restored the next time the application is started. Do not edit this file by hand.

Even though the firewall is activated immediately, the **iptables** service is not configured to start automatically at boot time. Refer to [Section 2.9.2.6, "Activating the IPTables Service"](#) for more information.

2.9.2.6. Activating the IPTables Service

The firewall rules are only active if the **iptables** service is running. To manually start the service, use the following command:

```
[root@myServer ~] # service iptables restart
```

To ensure that **iptables** starts when the system is booted, use the following command:

```
[root@myServer ~] # chkconfig --level 345 iptables on
```

2.9.3. Using IPTables

The first step in using **iptables** is to start the **iptables** service. Use the following command to start the **iptables** service:

```
[root@myServer ~] # service iptables start
```



Note

The **ip6tables** service can be turned off if you intend to use the **iptables** service only. If you deactivate the **ip6tables** service, remember to deactivate the IPv6 network also. Never leave a network device active without the matching firewall.

To force **iptables** to start by default when the system is booted, use the following command:

```
[root@myServer ~] # chkconfig --level 345 iptables on
```

This forces **iptables** to start whenever the system is booted into runlevel 3, 4, or 5.

2.9.3.1. IPTables Command Syntax

The following sample **iptables** command illustrates the basic command syntax:

```
[root@myServer ~ ] # iptables -A <chain> -j <target>
```

The `-A` option specifies that the rule be appended to `<chain>`. Each chain is comprised of one or more *rules*, and is therefore also known as a *ruleset*.

The three built-in chains are INPUT, OUTPUT, and FORWARD. These chains are permanent and cannot be deleted. The chain specifies the point at which a packet is manipulated.

The `-j <target>` option specifies the target of the rule; i.e., what to do if the packet matches the rule. Examples of built-in targets are ACCEPT, DROP, and REJECT.

Refer to the **iptables** man page for more information on the available chains, options, and targets.

2.9.3.2. Basic Firewall Policies

Establishing basic firewall policies creates a foundation for building more detailed, user-defined rules.

Each **iptables** chain is comprised of a default policy, and zero or more rules which work in concert with the default policy to define the overall ruleset for the firewall.

The default policy for a chain can be either DROP or ACCEPT. Security-minded administrators typically implement a default policy of DROP, and only allow specific packets on a case-by-case basis. For example, the following policies block all incoming and outgoing packets on a network gateway:

```
[root@myServer ~ ] # iptables -P INPUT DROP
[root@myServer ~ ] # iptables -P OUTPUT DROP
```

It is also recommended that any *forwarded packets* — network traffic that is to be routed from the firewall to its destination node — be denied as well, to restrict internal clients from inadvertent exposure to the Internet. To do this, use the following rule:

```
[root@myServer ~ ] # iptables -P FORWARD DROP
```

When you have established the default policies for each chain, you can create and save further rules for your particular network and security requirements.

The following sections describe how to save iptables rules and outline some of the rules you might implement in the course of building your iptables firewall.

2.9.3.3. Saving and Restoring IPTables Rules

Changes to **iptables** are transitory; if the system is rebooted or if the **iptables** service is restarted, the rules are automatically flushed and reset. To save the rules so that they are loaded when the **iptables** service is started, use the following command:

```
[root@myServer ~ ] # service iptables save
```

The rules are stored in the file `/etc/sysconfig/iptables` and are applied whenever the service is started or the machine is rebooted.

2.9.4. Common IPTables Filtering

Preventing remote attackers from accessing a LAN is one of the most important aspects of network security. The integrity of a LAN should be protected from malicious remote users through the use of stringent firewall rules.

However, with a default policy set to block all incoming, outgoing, and forwarded packets, it is impossible for the firewall/gateway and internal LAN users to communicate with each other or with external resources.

To allow users to perform network-related functions and to use networking applications, administrators must open certain ports for communication.

For example, to allow access to port 80 *on the firewall*, append the following rule:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp -m tcp --dport 80 -j ACCEPT
```

This allows users to browse websites that communicate using the standard port 80. To allow access to secure websites (for example, <https://www.example.com/>), you also need to provide access to port 443, as follows:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp -m tcp --dport 443 -j ACCEPT
```



Important

When creating an **iptables** ruleset, order is important.

If a rule specifies that any packets from the 192.168.100.0/24 subnet be dropped, and this is followed by a rule that allows packets from 192.168.100.13 (which is within the dropped subnet), then the second rule is ignored.

The rule to allow packets from 192.168.100.13 must precede the rule that drops the remainder of the subnet.

To insert a rule in a specific location in an existing chain, use the **-I** option. For example:

```
[root@myServer ~ ] # iptables -I INPUT 1 -i lo -p all -j ACCEPT
```

This rule is inserted as the first rule in the INPUT chain to allow local loopback device traffic.

There may be times when you require remote access to the LAN. Secure services, for example SSH, can be used for encrypted remote connection to LAN services.

Administrators with PPP-based resources (such as modem banks or bulk ISP accounts), dial-up access can be used to securely circumvent firewall barriers. Because they are direct connections, modem connections are typically behind a firewall/gateway.

For remote users with broadband connections, however, special cases can be made. You can configure **iptables** to accept connections from remote SSH clients. For example, the following rules allow remote SSH access:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp --dport 22 -j ACCEPT
```

```
[root@myServer ~ ] # iptables -A OUTPUT -p tcp --sport 22 -j ACCEPT
```

These rules allow incoming and outbound access for an individual system, such as a single PC directly connected to the Internet or a firewall/gateway. However, they do not allow nodes behind the firewall/gateway to access these services. To allow LAN access to these services, you can use *Network Address Translation (NAT)* with **iptables** filtering rules.

2.9.5. FORWARD and NAT Rules

Most ISPs provide only a limited number of publicly routable IP addresses to the organizations they serve.

Administrators must, therefore, find alternative ways to share access to Internet services without giving public IP addresses to every node on the LAN. Using private IP addresses is the most common way of allowing all nodes on a LAN to properly access internal and external network services.

Edge routers (such as firewalls) can receive incoming transmissions from the Internet and route the packets to the intended LAN node. At the same time, firewalls/gateways can also route outgoing requests from a LAN node to the remote Internet service.

This forwarding of network traffic can become dangerous at times, especially with the availability of modern cracking tools that can spoof *internal* IP addresses and make the remote attacker's machine act as a node on your LAN.

To prevent this, **iptables** provides routing and forwarding policies that can be implemented to prevent abnormal usage of network resources.

The FORWARD chain allows an administrator to control where packets can be routed within a LAN. For example, to allow forwarding for the entire LAN (assuming the firewall/gateway is assigned an internal IP address on eth1), use the following rules:

```
[root@myServer ~ ] # iptables -A FORWARD -i eth1 -j ACCEPT
[root@myServer ~ ] # iptables -A FORWARD -o eth1 -j ACCEPT
```

This rule gives systems behind the firewall/gateway access to the internal network. The gateway routes packets from one LAN node to its intended destination node, passing all packets through its **eth1** device.

Note

By default, the IPv4 policy in Fedora kernels disables support for IP forwarding. This prevents machines that run Fedora from functioning as dedicated edge routers. To enable IP forwarding, use the following command:

```
[root@myServer ~ ] # sysctl -w net.ipv4.ip_forward=1
```

This configuration change is only valid for the current session; it does not persist beyond a reboot or network service restart. To permanently set IP forwarding, edit the **/etc/sysctl.conf** file as follows:

Locate the following line:

```
net.ipv4.ip_forward = 0
```

Edit it to read as follows:

```
net.ipv4.ip_forward = 1
```

Use the following command to enable the change to the `sysctl.conf` file:

```
[root@myServer ~ ] # sysctl -p /etc/sysctl.conf
```

2.9.5.1. Postrouting and IP Masquerading

Accepting forwarded packets via the firewall's internal IP device allows LAN nodes to communicate with each other; however they still cannot communicate externally to the Internet.

To allow LAN nodes with private IP addresses to communicate with external public networks, configure the firewall for *IP masquerading*, which masks requests from LAN nodes with the IP address of the firewall's external device (in this case, `eth0`):

```
[root@myServer ~ ] # iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

This rule uses the NAT packet matching table (`-t nat`) and specifies the built-in `POSTROUTING` chain for NAT (`-A POSTROUTING`) on the firewall's external networking device (`-o eth0`).

`POSTROUTING` allows packets to be altered as they are leaving the firewall's external device.

The `-j MASQUERADE` target is specified to mask the private IP address of a node with the external IP address of the firewall/gateway.

2.9.5.2. Prerouting

If you have a server on your internal network that you want make available externally, you can use the `-j DNAT` target of the `PREROUTING` chain in NAT to specify a destination IP address and port where incoming packets requesting a connection to your internal service can be forwarded.

For example, if you want to forward incoming HTTP requests to your dedicated Apache HTTP Server at `172.31.0.23`, use the following command:

```
[root@myServer ~ ] # iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j DNAT --to 172.31.0.23:80
```

This rule specifies that the `nat` table use the built-in `PREROUTING` chain to forward incoming HTTP requests exclusively to the listed destination IP address of `172.31.0.23`.



Note

If you have a default policy of `DROP` in your `FORWARD` chain, you must append a rule to forward all incoming HTTP requests so that destination NAT routing is possible. To do this, use the following command:

```
[root@myServer ~ ] # iptables -A FORWARD -i eth0 -p tcp --dport 80 -d 172.31.0.23 -j ACCEPT
```

This rule forwards all incoming HTTP requests from the firewall to the intended destination; the Apache HTTP Server behind the firewall.

2.9.5.3. DMZs and IPTables

You can create **iptables** rules to route traffic to certain machines, such as a dedicated HTTP or FTP server, in a *demilitarized zone* (DMZ). A DMZ is a special local subnetwork dedicated to providing services on a public carrier, such as the Internet.

For example, to set a rule for routing incoming HTTP requests to a dedicated HTTP server at 10.0.4.2 (outside of the 192.168.1.0/24 range of the LAN), NAT uses the PREROUTING table to forward the packets to the appropriate destination:

```
[root@myServer ~ ] # iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j DNAT --to-destination 10.0.4.2:80
```

With this command, all HTTP connections to port 80 from outside of the LAN are routed to the HTTP server on a network separate from the rest of the internal network. This form of network segmentation can prove safer than allowing HTTP connections to a machine on the network.

If the HTTP server is configured to accept secure connections, then port 443 must be forwarded as well.

2.9.6. Malicious Software and Spoofed IP Addresses

More elaborate rules can be created that control access to specific subnets, or even specific nodes, within a LAN. You can also restrict certain dubious applications or programs such as trojans, worms, and other client/server viruses from contacting their server.

For example, some trojans scan networks for services on ports from 31337 to 31340 (called the *elite* ports in cracking terminology).

Since there are no legitimate services that communicate via these non-standard ports, blocking them can effectively diminish the chances that potentially infected nodes on your network independently communicate with their remote master servers.

The following rules drop all TCP traffic that attempts to use port 31337:

```
[root@myServer ~ ] # iptables -A OUTPUT -o eth0 -p tcp --dport 31337 --sport 31337 -j DROP
[root@myServer ~ ] # iptables -A FORWARD -o eth0 -p tcp --dport 31337 --sport 31337 -j DROP
```

You can also block outside connections that attempt to spoof private IP address ranges to infiltrate your LAN.

For example, if your LAN uses the 192.168.1.0/24 range, you can design a rule that instructs the Internet-facing network device (for example, eth0) to drop any packets to that device with an address in your LAN IP range.

Because it is recommended to reject forwarded packets as a default policy, any other spoofed IP address to the external-facing device (eth0) is rejected automatically.

```
[root@myServer ~ ] # iptables -A FORWARD -s 192.168.1.0/24 -i eth0 -j DROP
```



Note

There is a distinction between the DROP and REJECT targets when dealing with *appended* rules.

The REJECT target denies access and returns a `connection refused` error to users who attempt to connect to the service. The DROP target, as the name implies, drops the packet without any warning.

Administrators can use their own discretion when using these targets. However, to avoid user confusion and attempts to continue connecting, the REJECT target is recommended.

2.9.7. IPTables and Connection Tracking

You can inspect and restrict connections to services based on their *connection state*. A module within **iptables** uses a method called *connection tracking* to store information about incoming connections. You can allow or deny access based on the following connection states:

- NEW — A packet requesting a new connection, such as an HTTP request.
- ESTABLISHED — A packet that is part of an existing connection.
- RELATED — A packet that is requesting a new connection but is part of an existing connection. For example, FTP uses port 21 to establish a connection, but data is transferred on a different port (typically port 20).
- INVALID — A packet that is not part of any connections in the connection tracking table.

You can use the stateful functionality of **iptables** connection tracking with any network protocol, even if the protocol itself is stateless (such as UDP). The following example shows a rule that uses connection tracking to forward only the packets that are associated with an established connection:

```
[root@myServer ~ ] # iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT
```

2.9.8. IPv6

The introduction of the next-generation Internet Protocol, called IPv6, expands beyond the 32-bit address limit of IPv4 (or IP). IPv6 supports 128-bit addresses, and carrier networks that are IPv6 aware are therefore able to address a larger number of routable addresses than IPv4.

Fedora supports IPv6 firewall rules using the Netfilter 6 subsystem and the **ip6tables** command. In Fedora 12, both IPv4 and IPv6 services are enabled by default.

The **ip6tables** command syntax is identical to **iptables** in every aspect except that it supports 128-bit addresses. For example, use the following command to enable SSH connections on an IPv6-aware network server:

```
[root@myServer ~ ] # ip6tables -A INPUT -i eth0 -p tcp -s 3ffe:ffff:100::1/128 --dport 22 -j ACCEPT
```

For more information about IPv6 networking, refer to the IPv6 Information Page at <http://www.ipv6.org/>.

2.9.9. Additional Resources

There are several aspects to firewalls and the Linux Netfilter subsystem that could not be covered in this chapter. For more information, refer to the following resources.

2.9.9.1. Installed Firewall Documentation

- Refer to [Section 2.10, “IPTables”](#) for more detailed information on the **iptables** command, including definitions for many command options.
- The **iptables** man page contains a brief summary of the various options.

2.9.9.2. Useful Firewall Websites

- <http://www.netfilter.org/> — The official homepage of the Netfilter and **iptables** project.
- <http://www.tldp.org/> — The Linux Documentation Project contains several useful guides relating to firewall creation and administration.
- <http://www.iana.org/assignments/port-numbers> — The official list of registered and common service ports as assigned by the Internet Assigned Numbers Authority.

2.9.9.3. Related Documentation

- *Red Hat Linux Firewalls*, by Bill McCarty; Red Hat Press — a comprehensive reference to building network and server firewalls using open source packet filtering technology such as Netfilter and **iptables**. It includes topics that cover analyzing firewall logs, developing firewall rules, and customizing your firewall using various graphical tools.
- *Linux Firewalls*, by Robert Ziegler; New Riders Press — contains a wealth of information on building firewalls using both 2.2 kernel **ipchains** as well as Netfilter and **iptables**. Additional security topics such as remote access issues and intrusion detection systems are also covered.

2.10. IPTables

Included with Fedora are advanced tools for network *packet filtering* — the process of controlling network packets as they enter, move through, and exit the network stack within the kernel. Kernel versions prior to 2.4 relied on **ipchains** for packet filtering and used lists of rules applied to packets at each step of the filtering process. The 2.4 kernel introduced **iptables** (also called *netfilter*), which is similar to **ipchains** but greatly expands the scope and control available for filtering network packets.

This chapter focuses on packet filtering basics, explains various options available with **iptables** commands, and explains how filtering rules can be preserved between system reboots.

Refer to [Section 2.10.6, “Additional Resources”](#) for instructions on how to construct **iptables** rules and setting up a firewall based on these rules.



Important

The default firewall mechanism in the 2.4 and later kernels is **iptables**, but **iptables** cannot be used if **ipchains** is already running. If **ipchains** is present at boot time, the kernel issues an error and fails to start **iptables**.

The functionality of `ipchains` is not affected by these errors.

2.10.1. Packet Filtering

The Linux kernel uses the **Netfilter** facility to filter packets, allowing some of them to be received by or pass through the system while stopping others. This facility is built in to the Linux kernel, and has three built-in *tables* or *rules lists*, as follows:

- `filter` — The default table for handling network packets.
- `nat` — Used to alter packets that create a new connection and used for *Network Address Translation (NAT)*.
- `mangle` — Used for specific types of packet alteration.

Each table has a group of built-in *chains*, which correspond to the actions performed on the packet by **netfilter**.

The built-in chains for the `filter` table are as follows:

- `INPUT` — Applies to network packets that are targeted for the host.
- `OUTPUT` — Applies to locally-generated network packets.
- `FORWARD` — Applies to network packets routed through the host.

The built-in chains for the `nat` table are as follows:

- `PREROUTING` — Alters network packets when they arrive.
- `OUTPUT` — Alters locally-generated network packets before they are sent out.
- `POSTROUTING` — Alters network packets before they are sent out.

The built-in chains for the `mangle` table are as follows:

- `INPUT` — Alters network packets targeted for the host.
- `OUTPUT` — Alters locally-generated network packets before they are sent out.
- `FORWARD` — Alters network packets routed through the host.
- `PREROUTING` — Alters incoming network packets before they are routed.
- `POSTROUTING` — Alters network packets before they are sent out.

Every network packet received by or sent from a Linux system is subject to at least one table. However, a packet may be subjected to multiple rules within each table before emerging at the end of the chain. The structure and purpose of these rules may vary, but they usually seek to identify a packet coming from or going to a particular IP address, or set of addresses, when using a particular protocol and network service.



Note

By default, firewall rules are saved in the `/etc/sysconfig/iptables` or `/etc/sysconfig/ip6tables` files.

The `iptables` service starts before any DNS-related services when a Linux system is booted. This means that firewall rules can only reference numeric IP addresses (for example, `192.168.0.1`). Domain names (for example, `host.example.com`) in such rules produce errors.

Regardless of their destination, when packets match a particular rule in one of the tables, a *target* or action is applied to them. If the rule specifies an **ACCEPT** target for a matching packet, the packet skips the rest of the rule checks and is allowed to continue to its destination. If a rule specifies a **DROP** target, that packet is refused access to the system and nothing is sent back to the host that sent the packet. If a rule specifies a **QUEUE** target, the packet is passed to user-space. If a rule specifies the optional **REJECT** target, the packet is dropped, but an error packet is sent to the packet's originator.

Every chain has a default policy to **ACCEPT**, **DROP**, **REJECT**, or **QUEUE**. If none of the rules in the chain apply to the packet, then the packet is dealt with in accordance with the default policy.

The `iptables` command configures these tables, as well as sets up new tables if necessary.

2.10.2. Command Options for IPTables

Rules for filtering packets are created using the `iptables` command. The following aspects of the packet are most often used as criteria:

- *Packet Type* — Specifies the type of packets the command filters.
- *Packet Source/Destination* — Specifies which packets the command filters based on the source or destination of the packet.
- *Target* — Specifies what action is taken on packets matching the above criteria.

Refer to [Section 2.10.2.4, “IPTables Match Options”](#) and [Section 2.10.2.5, “Target Options”](#) for more information about specific options that address these aspects of a packet.

The options used with specific `iptables` rules must be grouped logically, based on the purpose and conditions of the overall rule, for the rule to be valid. The remainder of this section explains commonly-used options for the `iptables` command.

2.10.2.1. Structure of IPTables Command Options

Many `iptables` commands have the following structure:

```
iptables [-t <table-name>] <command> <chain-name> \ <parameter-1> <option-1> \ <parameter-n> <option-n>
```

`<table-name>` — Specifies which table the rule applies to. If omitted, the `filter` table is used.

`<command>` — Specifies the action to perform, such as appending or deleting a rule.

`<chain-name>` — Specifies the chain to edit, create, or delete.

<parameter>-<option> pairs — Parameters and associated options that specify how to process a packet that matches the rule.

The length and complexity of an **iptables** command can change significantly, based on its purpose.

For example, a command to remove a rule from a chain can be very short:

```
iptables -D <chain-name> <line-number>
```

In contrast, a command that adds a rule which filters packets from a particular subnet using a variety of specific parameters and options can be rather long. When constructing **iptables** commands, it is important to remember that some parameters and options require further parameters and options to construct a valid rule. This can produce a cascading effect, with the further parameters requiring yet more parameters. Until every parameter and option that requires another set of options is satisfied, the rule is not valid.

Type **iptables -h** to view a comprehensive list of **iptables** command structures.

2.10.2.2. Command Options

Command options instruct **iptables** to perform a specific action. Only one command option is allowed per **iptables** command. With the exception of the help command, all commands are written in upper-case characters.

The **iptables** commands are as follows:

- **-A** — Appends the rule to the end of the specified chain. Unlike the **-I** option described below, it does not take an integer argument. It always appends the rule to the end of the specified chain.
- **-C** — Checks a particular rule before adding it to the user-specified chain. This command can help you construct complex **iptables** rules by prompting you for additional parameters and options.
- **-D <integer> | <rule>** — Deletes a rule in a particular chain by number (such as 5 for the fifth rule in a chain), or by rule specification. The rule specification must exactly match an existing rule.
- **-E** — Renames a user-defined chain. A user-defined chain is any chain other than the default, pre-existing chains. (Refer to the **-N** option, below, for information on creating user-defined chains.) This is a cosmetic change and does not affect the structure of the table.



Note

If you attempt to rename one of the default chains, the system reports a `Match not found` error. You cannot rename the default chains.

- **-F** — Flushes the selected chain, which effectively deletes every rule in the chain. If no chain is specified, this command flushes every rule from every chain.
- **-h** — Provides a list of command structures, as well as a quick summary of command parameters and options.
- **-I [<integer>]** — Inserts the rule in the specified chain at a point specified by a user-defined integer argument. If no argument is specified, the rule is inserted at the top of the chain.



Important

As noted above, the order of rules in a chain determines which rules apply to which packets. This is important to remember when adding rules using either the `-A` or `-I` option.

This is especially important when adding rules using the `-I` with an integer argument. If you specify an existing number when adding a rule to a chain, **iptables** adds the new rule *before* (or above) the existing rule.

- `-L` — Lists all of the rules in the chain specified after the command. To list all rules in all chains in the default `filter` table, do not specify a chain or table. Otherwise, the following syntax should be used to list the rules in a specific chain in a particular table:

```
iptables -L <chain-name> -t <table-name>
```

Additional options for the `-L` command option, which provide rule numbers and allow more verbose rule descriptions, are described in [Section 2.10.2.6, “Listing Options”](#).

- `-N` — Creates a new chain with a user-specified name. The chain name must be unique, otherwise an error message is displayed.
- `-P` — Sets the default policy for the specified chain, so that when packets traverse an entire chain without matching a rule, they are sent to the specified target, such as `ACCEPT` or `DROP`.
- `-R` — Replaces a rule in the specified chain. The rule's number must be specified after the chain's name. The first rule in a chain corresponds to rule number one.
- `-X` — Deletes a user-specified chain. You cannot delete a built-in chain.
- `-Z` — Sets the byte and packet counters in all chains for a table to zero.

2.10.2.3. IPTables Parameter Options

Certain **iptables** commands, including those used to add, append, delete, insert, or replace rules within a particular chain, require various parameters to construct a packet filtering rule.

- `-c` — Resets the counters for a particular rule. This parameter accepts the `PKTS` and `BYTES` options to specify which counter to reset.
- `-d` — Sets the destination hostname, IP address, or network of a packet that matches the rule. When matching a network, the following IP address/netmask formats are supported:
 - `N.N.N.N/M.M.M.M` — Where `N.N.N.N` is the IP address range and `M.M.M.M` is the netmask.
 - `N.N.N.N/M` — Where `N.N.N.N` is the IP address range and `M` is the bitmask.
- `-f` — Applies this rule only to fragmented packets.

You can use the exclamation point character (!) option after this parameter to specify that only unfragmented packets are matched.



Note

Distinguishing between fragmented and unfragmented packets is desirable, despite fragmented packets being a standard part of the IP protocol.

Originally designed to allow IP packets to travel over networks with differing frame sizes, these days fragmentation is more commonly used to generate DoS attacks using mal-formed packets. It's also worth noting that IPv6 disallows fragmentation entirely.

- `-i` — Sets the incoming network interface, such as `eth0` or `ppp0`. With **iptables**, this optional parameter may only be used with the INPUT and FORWARD chains when used with the `filter` table and the PREROUTING chain with the `nat` and `mangle` tables.

This parameter also supports the following special options:

- Exclamation point character (!) — Reverses the directive, meaning any specified interfaces are excluded from this rule.
- Plus character (+) — A wildcard character used to match all interfaces that match the specified string. For example, the parameter `-i eth+` would apply this rule to any Ethernet interfaces but exclude any other interfaces, such as `ppp0`.

If the `-i` parameter is used but no interface is specified, then every interface is affected by the rule.

- `-j` — Jumps to the specified target when a packet matches a particular rule.

The standard targets are ACCEPT, DROP, QUEUE, and RETURN.

Extended options are also available through modules loaded by default with the Fedora **iptables** RPM package. Valid targets in these modules include LOG, MARK, and REJECT, among others. Refer to the **iptables** man page for more information about these and other targets.

This option can also be used to direct a packet matching a particular rule to a user-defined chain outside of the current chain so that other rules can be applied to the packet.

If no target is specified, the packet moves past the rule with no action taken. The counter for this rule, however, increases by one.

- `-o` — Sets the outgoing network interface for a rule. This option is only valid for the OUTPUT and FORWARD chains in the `filter` table, and the POSTROUTING chain in the `nat` and `mangle` tables. This parameter accepts the same options as the incoming network interface parameter (`-i`).
- `-p <protocol>` — Sets the IP protocol affected by the rule. This can be either `icmp`, `tcp`, `udp`, or `all`, or it can be a numeric value, representing one of these or a different protocol. You can also use any protocols listed in the `/etc/protocols` file.

The "all" protocol means the rule applies to every supported protocol. If no protocol is listed with this rule, it defaults to "all".

- `-s` — Sets the source for a particular packet using the same syntax as the destination (`-d`) parameter.

2.10.2.4. IPTables Match Options

Different network protocols provide specialized matching options which can be configured to match a particular packet using that protocol. However, the protocol must first be specified in the **iptables** command. For example, `-p <protocol-name>` enables options for the specified protocol. Note that you can also use the protocol ID, instead of the protocol name. Refer to the following examples, each of which have the same effect:

```
iptables -A INPUT -p icmp --icmp-type any -j ACCEPT
```

```
iptables -A INPUT -p 5813 --icmp-type any -j ACCEPT
```

Service definitions are provided in the `/etc/services` file. For readability, it is recommended that you use the service names rather than the port numbers.



Warning

Secure the `/etc/services` file to prevent unauthorized editing. If this file is editable, crackers can use it to enable ports on your machine you have otherwise closed. To secure this file, type the following commands as root:

```
[root@myServer ~]# chown root.root /etc/services
[root@myServer ~]# chmod 0644 /etc/services
[root@myServer ~]# chattr +i /etc/services
```

This prevents the file from being renamed, deleted or having links made to it.

2.10.2.4.1. TCP Protocol

These match options are available for the TCP protocol (`-p tcp`):

- `--dport` — Sets the destination port for the packet.

To configure this option, use a network service name (such as `www` or `smtp`); a port number; or a range of port numbers.

To specify a range of port numbers, separate the two numbers with a colon (`:`). For example: `-p tcp --dport 3000:3200`. The largest acceptable valid range is `0:65535`.

Use an exclamation point character (`!`) after the `--dport` option to match all packets that *do not* use that network service or port.

To browse the names and aliases of network services and the port numbers they use, view the `/etc/services` file.

The `--destination-port` match option is synonymous with `--dport`.

- `--sport` — Sets the source port of the packet using the same options as `--dport`. The `--source-port` match option is synonymous with `--sport`.
- `--syn` — Applies to all TCP packets designed to initiate communication, commonly called *SYN packets*. Any packets that carry a data payload are not touched.

Use an exclamation point character (!) after the `--syn` option to match all non-SYN packets.

- `--tcp-flags <tested flag list> <set flag list>` — Allows TCP packets that have specific bits (flags) set, to match a rule.

The `--tcp-flags` match option accepts two parameters. The first parameter is the mask; a comma-separated list of flags to be examined in the packet. The second parameter is a comma-separated list of flags that must be set for the rule to match.

The possible flags are:

- ACK
- FIN
- PSH
- RST
- SYN
- URG
- ALL
- NONE

For example, an **iptables** rule that contains the following specification only matches TCP packets that have the SYN flag set and the ACK and FIN flags not set:

```
--tcp-flags ACK,FIN,SYN SYN
```

Use the exclamation point character (!) after the `--tcp-flags` to reverse the effect of the match option.

- `--tcp-option` — Attempts to match with TCP-specific options that can be set within a particular packet. This match option can also be reversed with the exclamation point character (!).

2.10.2.4.2. UDP Protocol

These match options are available for the UDP protocol (`-p udp`):

- `--dport` — Specifies the destination port of the UDP packet, using the service name, port number, or range of port numbers. The `--destination-port` match option is synonymous with `--dport`.
- `--sport` — Specifies the source port of the UDP packet, using the service name, port number, or range of port numbers. The `--source-port` match option is synonymous with `--sport`.

For the `--dport` and `--sport` options, to specify a range of port numbers, separate the two numbers with a colon (:). For example: `-p tcp --dport 3000:3200`. The largest acceptable valid range is 0:65535.

2.10.2.4.3. ICMP Protocol

The following match options are available for the Internet Control Message Protocol (ICMP) (-p icmp):

- `--icmp-type` — Sets the name or number of the ICMP type to match with the rule. A list of valid ICMP names can be retrieved by typing the `iptables -p icmp -h` command.

2.10.2.4.4. Additional Match Option Modules

Additional match options are available through modules loaded by the `iptables` command.

To use a match option module, load the module by name using the `-m <module-name>`, where `<module-name>` is the name of the module.

Many modules are available by default. You can also create modules to provide additional functionality.

The following is a partial list of the most commonly used modules:

- `limit` module — Places limits on how many packets are matched to a particular rule.

When used in conjunction with the **LOG** target, the `limit` module can prevent a flood of matching packets from filling up the system log with repetitive messages or using up system resources.

Refer to [Section 2.10.2.5, “Target Options”](#) for more information about the **LOG** target.

The `limit` module enables the following options:

- `--limit` — Sets the maximum number of matches for a particular time period, specified as a `<value>/<period>` pair. For example, using `--limit 5/hour` allows five rule matches per hour.

Periods can be specified in seconds, minutes, hours, or days.

If a number and time modifier are not used, the default value of 3/hour is assumed.

- `--limit-burst` — Sets a limit on the number of packets able to match a rule at one time.

This option is specified as an integer and should be used in conjunction with the `--limit` option.

If no value is specified, the default value of five (5) is assumed.

- `state` module — Enables state matching.

The `state` module enables the following options:

- `--state` — match a packet with the following connection states:
 - **ESTABLISHED** — The matching packet is associated with other packets in an established connection. You need to accept this state if you want to maintain a connection between a client and a server.
 - **INVALID** — The matching packet cannot be tied to a known connection.
 - **NEW** — The matching packet is either creating a new connection or is part of a two-way connection not previously seen. You need to accept this state if you want to allow new connections to a service.

- **RELATED** — The matching packet is starting a new connection related in some way to an existing connection. An example of this is FTP, which uses one connection for control traffic (port 21), and a separate connection for data transfer (port 20).

These connection states can be used in combination with one another by separating them with commas, such as `-m state --state INVALID,NEW`.

- **mac module** — Enables hardware MAC address matching.

The mac module enables the following option:

- `--mac-source` — Matches a MAC address of the network interface card that sent the packet. To exclude a MAC address from a rule, place an exclamation point character (!) after the `--mac-source` match option.

Refer to the **iptables** man page for more match options available through modules.

2.10.2.5. Target Options

When a packet has matched a particular rule, the rule can direct the packet to a number of different targets which determine the appropriate action. Each chain has a default target, which is used if none of the rules on that chain match a packet or if none of the rules which match the packet specify a target.

The following are the standard targets:

- `<user-defined-chain>` — A user-defined chain within the table. User-defined chain names must be unique. This target passes the packet to the specified chain.
- **ACCEPT** — Allows the packet through to its destination or to another chain.
- **DROP** — Drops the packet without responding to the requester. The system that sent the packet is not notified of the failure.
- **QUEUE** — The packet is queued for handling by a user-space application.
- **RETURN** — Stops checking the packet against rules in the current chain. If the packet with a RETURN target matches a rule in a chain called from another chain, the packet is returned to the first chain to resume rule checking where it left off. If the RETURN rule is used on a built-in chain and the packet cannot move up to its previous chain, the default target for the current chain is used.

In addition, extensions are available which allow other targets to be specified. These extensions are called target modules or match option modules and most only apply to specific tables and situations. Refer to [Section 2.10.2.4.4, “Additional Match Option Modules”](#) for more information about match option modules.

Many extended target modules exist, most of which only apply to specific tables or situations. Some of the most popular target modules included by default in Fedora are:

- **LOG** — Logs all packets that match this rule. Because the packets are logged by the kernel, the `/etc/syslog.conf` file determines where these log entries are written. By default, they are placed in the `/var/log/messages` file.

Additional options can be used after the LOG target to specify the way in which logging occurs:

- `--log-level` — Sets the priority level of a logging event. Refer to the `syslog.conf` man page for a list of priority levels.
- `--log-ip-options` — Logs any options set in the header of an IP packet.
- `--log-prefix` — Places a string of up to 29 characters before the log line when it is written. This is useful for writing syslog filters for use in conjunction with packet logging.



Note

Due to an issue with this option, you should add a trailing space to the `log-prefix` value.

- `--log-tcp-options` — Logs any options set in the header of a TCP packet.
- `--log-tcp-sequence` — Writes the TCP sequence number for the packet in the log.
- REJECT — Sends an error packet back to the remote system and drops the packet.

The REJECT target accepts `--reject-with <type>` (where `<type>` is the rejection type) allowing more detailed information to be returned with the error packet. The message `port-unreachable` is the default error type given if no other option is used. Refer to the `iptables` man page for a full list of `<type>` options.

Other target extensions, including several that are useful for IP masquerading using the `nat` table, or with packet alteration using the `mangle` table, can be found in the `iptables` man page.

2.10.2.6. Listing Options

The default list command, `iptables -L [<chain-name>]`, provides a very basic overview of the default filter table's current chains. Additional options provide more information:

- `-v` — Displays verbose output, such as the number of packets and bytes each chain has processed, the number of packets and bytes each rule has matched, and which interfaces apply to a particular rule.
- `-x` — Expands numbers into their exact values. On a busy system, the number of packets and bytes processed by a particular chain or rule may be abbreviated to Kilobytes, Megabytes (Megabytes) or Gigabytes. This option forces the full number to be displayed.
- `-n` — Displays IP addresses and port numbers in numeric format, rather than the default hostname and network service format.
- `--line-numbers` — Lists rules in each chain next to their numeric order in the chain. This option is useful when attempting to delete the specific rule in a chain or to locate where to insert a rule within a chain.
- `-t <table-name>` — Specifies a table name. If omitted, defaults to the filter table.

2.10.3. Saving IPTables Rules

Rules created with the **iptables** command are stored in memory. If the system is restarted before saving the **iptables** rule set, all rules are lost. For netfilter rules to persist through a system reboot, they need to be saved. To save netfilter rules, type the following command as root:

```
/sbin/service iptables save
```

This executes the **iptables** init script, which runs the **/sbin/iptables-save** program and writes the current **iptables** configuration to **/etc/sysconfig/iptables**. The existing **/etc/sysconfig/iptables** file is saved as **/etc/sysconfig/iptables.save**.

The next time the system boots, the **iptables** init script reapplies the rules saved in **/etc/sysconfig/iptables** by using the **/sbin/iptables-restore** command.

While it is always a good idea to test a new **iptables** rule before committing it to the **/etc/sysconfig/iptables** file, it is possible to copy **iptables** rules into this file from another system's version of this file. This provides a quick way to distribute sets of **iptables** rules to multiple machines.

You can also save the **iptables** rules to a separate file for distribution, backup or other purposes. To save your **iptables** rules, type the following command as root:

```
[root@myServer ~]# iptables-save > <filename>where <filename> is a user-defined name for your ruleset.
```



Important

If distributing the **/etc/sysconfig/iptables** file to other machines, type **/sbin/service iptables restart** for the new rules to take effect.



Note

Note the difference between the **iptables** *command* (**/sbin/iptables**), which is used to manipulate the tables and chains that constitute the **iptables** functionality, and the **iptables** *service* (**/sbin/iptables service**), which is used to enable and disable the **iptables** service itself.

2.10.4. IPTables Control Scripts

There are two basic methods for controlling **iptables** in Fedora:

- **Firewall Configuration Tool (system-config-securitylevel)** — A graphical interface for creating, activating, and saving basic firewall rules. Refer to [Section 2.9.2, “Basic Firewall Configuration”](#) for more information.
- **/sbin/service iptables <option>** — Used to manipulate various functions of **iptables** using its initscript. The following options are available:
 - **start** — If a firewall is configured (that is, **/etc/sysconfig/iptables** exists), all running **iptables** are stopped completely and then started using the **/sbin/iptables-restore**

command. This option only works if the **ipchains** kernel module is not loaded. To check if this module is loaded, type the following command as root:

```
[root@MyServer ~]# lsmod | grep ipchains
```

If this command returns no output, it means the module is not loaded. If necessary, use the **/sbin/rmmod** command to remove the module.

- **stop** — If a firewall is running, the firewall rules in memory are flushed, and all iptables modules and helpers are unloaded.

If the **IPTABLES_SAVE_ON_STOP** directive in the **/etc/sysconfig/iptables-config** configuration file is changed from its default value to **yes**, current rules are saved to **/etc/sysconfig/iptables** and any existing rules are moved to the file **/etc/sysconfig/iptables.save**.

Refer to [Section 2.10.4.1, “IPTables Control Scripts Configuration File”](#) for more information about the **iptables-config** file.

- **restart** — If a firewall is running, the firewall rules in memory are flushed, and the firewall is started again if it is configured in **/etc/sysconfig/iptables**. This option only works if the **ipchains** kernel module is not loaded.

If the **IPTABLES_SAVE_ON_RESTART** directive in the **/etc/sysconfig/iptables-config** configuration file is changed from its default value to **yes**, current rules are saved to **/etc/sysconfig/iptables** and any existing rules are moved to the file **/etc/sysconfig/iptables.save**.

Refer to [Section 2.10.4.1, “IPTables Control Scripts Configuration File”](#) for more information about the **iptables-config** file.

- **status** — Displays the status of the firewall and lists all active rules.

The default configuration for this option displays IP addresses in each rule. To display domain and hostname information, edit the **/etc/sysconfig/iptables-config** file and change the value of **IPTABLES_STATUS_NUMERIC** to **no**. Refer to [Section 2.10.4.1, “IPTables Control Scripts Configuration File”](#) for more information about the **iptables-config** file.

- **panic** — Flushes all firewall rules. The policy of all configured tables is set to **DROP**.

This option could be useful if a server is known to be compromised. Rather than physically disconnecting from the network or shutting down the system, you can use this option to stop all further network traffic but leave the machine in a state ready for analysis or other forensics.

- **save** — Saves firewall rules to **/etc/sysconfig/iptables** using **iptables-save**. Refer to [Section 2.10.3, “Saving IPTables Rules”](#) for more information.

Note

To use the same initscript commands to control netfilter for IPv6, substitute **ip6tables** for **iptables** in the **/sbin/service** commands listed in this section. For more information about IPv6 and netfilter, refer to [Section 2.10.5, “IPTables and IPv6”](#).

2.10.4.1. IPTables Control Scripts Configuration File

The behavior of the **iptables** initscripts is controlled by the `/etc/sysconfig/iptables-config` configuration file. The following is a list of directives contained in this file:

- **IPTABLES_MODULES** — Specifies a space-separated list of additional **iptables** modules to load when a firewall is activated. These can include connection tracking and NAT helpers.
- **IPTABLES_MODULES_UNLOAD** — Unloads modules on restart and stop. This directive accepts the following values:
 - **yes** — The default value. This option must be set to achieve a correct state for a firewall restart or stop.
 - **no** — This option should only be set if there are problems unloading the netfilter modules.
- **IPTABLES_SAVE_ON_STOP** — Saves current firewall rules to `/etc/sysconfig/iptables` when the firewall is stopped. This directive accepts the following values:
 - **yes** — Saves existing rules to `/etc/sysconfig/iptables` when the firewall is stopped, moving the previous version to the `/etc/sysconfig/iptables.save` file.
 - **no** — The default value. Does not save existing rules when the firewall is stopped.
- **IPTABLES_SAVE_ON_RESTART** — Saves current firewall rules when the firewall is restarted. This directive accepts the following values:
 - **yes** — Saves existing rules to `/etc/sysconfig/iptables` when the firewall is restarted, moving the previous version to the `/etc/sysconfig/iptables.save` file.
 - **no** — The default value. Does not save existing rules when the firewall is restarted.
- **IPTABLES_SAVE_COUNTER** — Saves and restores all packet and byte counters in all chains and rules. This directive accepts the following values:
 - **yes** — Saves the counter values.
 - **no** — The default value. Does not save the counter values.
- **IPTABLES_STATUS_NUMERIC** — Outputs IP addresses in numeric form instead of domain or hostnames. This directive accepts the following values:
 - **yes** — The default value. Returns only IP addresses within a status output.
 - **no** — Returns domain or hostnames within a status output.

2.10.5. IPTables and IPv6

If the **iptables-ipv6** package is installed, netfilter in Fedora can filter the next-generation IPv6 Internet protocol. The command used to manipulate the IPv6 netfilter is **ip6tables**.

Most directives for this command are identical to those used for **iptables**, except the **nat** table is not yet supported. This means that it is not yet possible to perform IPv6 network address translation tasks, such as masquerading and port forwarding.

Rules for **ip6tables** are saved in the `/etc/sysconfig/ip6tables` file. Previous rules saved by the **ip6tables** initscripts are saved in the `/etc/sysconfig/ip6tables.save` file.

Configuration options for the **ip6tables** init script are stored in `/etc/sysconfig/ip6tables-config`, and the names for each directive vary slightly from their **iptables** counterparts.

For example, the **iptables-config** directive **IPTABLES_MODULES**:the equivalent in the **ip6tables-config** file is **IP6TABLES_MODULES**.

2.10.6. Additional Resources

Refer to the following sources for additional information on packet filtering with **iptables**.

- [Section 2.9, “Firewalls”](#) — Contains a chapter about the role of firewalls within an overall security strategy as well as strategies for constructing firewall rules.

2.10.6.1. Installed IP Tables Documentation

- **man iptables** — Contains a description of **iptables** as well as a comprehensive list of targets, options, and match extensions.

2.10.6.2. Useful IP Tables Websites

- <http://www.netfilter.org/> — The home of the netfilter/iptables project. Contains assorted information about **iptables**, including a FAQ addressing specific problems and various helpful guides by Rusty Russell, the Linux IP firewall maintainer. The HOWTO documents on the site cover subjects such as basic networking concepts, kernel packet filtering, and NAT configurations.
- http://www.linuxnewbie.org/nhf/Security/IPtables_Basics.html — An introduction to the way packets move through the Linux kernel, plus an introduction to constructing basic **iptables** commands.

Encryption

There are two main types of data that must be protected: data at rest and data in motion. These different types of data are protected in similar ways using similar technology but the implementations can be completely different. No single protective implementation can prevent all possible methods of compromise as the same information may be at rest and in motion at different points in time.

3.1. Data at Rest

Data at rest is data that is stored on a hard drive, tape, CD, DVD, disk, or other media. This information's biggest threat comes from being physically stolen. Laptops in airports, CDs going through the mail, and backup tapes that get left in the wrong places are all examples of events where data can be compromised through theft. If the data was encrypted on the media then you wouldn't have to worry as much about the data being compromised.

3.2. Full Disk Encryption

Full disk or partition encryption is one of the best ways of protecting your data. Not only is each file protected but also the temporary storage that may contain parts of these files is also protected. Full disk encryption will protect all of your files so you don't have to worry about selecting what you want to protect and possibly missing a file.

Fedora 9, and later, natively supports LUKS Encryption. LUKS will bulk encrypt your hard drive partitions so that while your computer is off your data is protected. This will also protect your computer from attackers attempting to use single-user-mode to login to your computer or otherwise gain access.

Full disk encryption solutions like LUKS only protect the data when your computer is off. Once the computer is on and LUKS has decrypted the disk, the files on that disk are available to anyone who would normally have access to them. To protect your files when the computer is on, use full disk encryption in combination with another solution such as file based encryption. Also remember to lock your computer whenever you are away from it. A passphrase protected screen saver set to activate after a few minutes of inactivity is a good way to keep intruders out.

3.3. File Based Encryption

GnuPG (GPG) is an open source version of PGP that allows you to sign and/or encrypt a file or an email message. This is useful to maintain integrity of the message or file and also protects the confidentiality of the information contained within the file or email. In the case of email, GPG provides dual protection. Not only can it provide Data at Rest protection but also Data In Motion protection once the message has been sent across the network.

File based encryption is intended to protect a file after it has left your computer, such as when you send a CD through the mail. Some file based encryption solutions will leave remnants of the encrypted files that an attacker who has physical access to your computer can recover under some circumstances. To protect the contents of those files from attackers who may have access to your computer, use file based encryption combined with another solution such as full disk encryption.

3.4. Data in Motion

Data in motion is data that is being transmitted over a network. The biggest threats to data in motion are interception and alteration. Your user name and password should never be transmitted over a network without protection as it could be intercepted and used by someone else to impersonate you

or gain access to sensitive information. Other private information such as bank account information should also be protected when transmitted across a network. If the network session was encrypted then you would not have to worry as much about the data being compromised while it is being transmitted.

Data in motion is particularly vulnerable to attackers because the attacker does not have to be near the computer in which the data is being stored rather they only have to be somewhere along the path. Encryption tunnels can protect data along the path of communications.

3.5. Virtual Private Networks

Virtual Private Networks (VPN) provide encrypted tunnels between computers or networks of computers across all ports. With a VPN in place, all network traffic from the client is forwarded to the server through the encrypted tunnel. This means that the client is logically on the same network as the server it is connected to via the VPN. VPNs are very common and are simple to use and setup.

3.6. Secure Shell

Secure Shell (SSH) is a powerful network protocol used to communicate with another system over a secure channel. The transmissions over SSH are encrypted and protected from interception. Cryptographic log-on can also be utilized to provide a better authentication method over traditional usernames and passwords.

SSH is very easy to activate. By simply starting the `sshd` service, the system will begin to accept connections and will allow access to the system when a correct username and password is provided during the connection process. The standard TCP port for the SSH service is 22, however this can be changed by modifying the configuration file `/etc/ssh/sshd_config` and restarting the service. This file also contains other configuration options for SSH.

Secure Shell (SSH) also provides encrypted tunnels between computers but only using a single port. [Port forwarding can be done over an SSH tunnel¹](#) and traffic will be encrypted as it passes over that tunnel but using port forwarding is not as fluid as a VPN.

3.7. LUKS Disk Encryption

Linux Unified Key Setup-on-disk-format (or LUKS) allows you to encrypt partitions on your Linux computer. This is particularly important when it comes to mobile computers and removable media. LUKS allows multiple user keys to decrypt a master key which is used for the bulk encryption of the partition.

3.7.1. LUKS Implementation in Fedora

Fedora 9, and later, utilizes LUKS to perform file system encryption. By default, the option to encrypt the file system is unchecked during the installation. If you select the option to encrypt you hard drive, you will be prompted for a passphrase that will be asked every time you boot the computer. This passphrase "unlocks" the bulk encryption key that is used to decrypt your partition. If you choose to modify the default partition table you can choose which partitions you want to encrypt. This is set in the partition table settings

Fedora's default implementation of LUKS is AES 128 with a SHA256 hashing. Ciphers that are available are:

¹ <http://www.redhatmagazine.com/2007/11/27/advanced-ssh-configuration-and-tunneling-we-dont-need-no-stinking-vpn-software>

- AES - Advanced Encryption Standard - [FIPS PUB 197](#)²
- Twofish (A 128-bit Block Cipher)
- Serpent
- cast5 - [RFC 2144](#)³
- cast6 - [RFC 2612](#)⁴

3.7.2. Manually Encrypting Directories



Warning

Following this procedure will remove all data on the partition that you are encrypting. You WILL lose all your information! Make sure you backup your data to an external source before beginning this procedure!

If you are running a version of Fedora prior to Fedora 9 and want to encrypt a partition, or you want to encrypt a partition after the installation of the current version of Fedora, the following directions are for you. The below example demonstrates encrypting your /home partition but any partition can be used.

The following procedure will wipe all your existing data, so be sure to have a tested backup before you start. This also requires you to have a separate partition for /home (in my case that is /dev/VG00/LV_home). All the following must be done as root. Any of these steps failing means you must not continue until the step succeeded.

3.7.3. Step-by-Step Instructions

1. enter runlevel 1: `telinit 1`
2. unmount your existing /home: `umount /home`
3. if it fails use `fuser` to find and kill processes hogging /home: `fuser -mvk /home`
4. verify /home is not mounted any longer: `cat /proc/mounts | grep home`
5. Fill your partition with random data: `dd if=/dev/urandom of=/dev/VG00/LV_home` This process takes many hours to complete.



Important

The process, however, is imperative in order to have good protection against break-in attempts. Just let it run overnight.

6. initialize your partition: `cryptsetup --verbose --verify-passphrase luksFormat /dev/VG00/LV_home`
7. open the newly encrypted device: `cryptsetup luksOpen /dev/VG00/LV_home home`
8. check it's there: `ls -l /dev/mapper | grep home`

9. create a filesystem: `mkfs.ext3 /dev/mapper/home`
10. mount it: `mount /dev/mapper/home /home`
11. check it's visible: `df -h | grep home`
12. add the following to `/etc/crypttab`: `home /dev/VG00/LV_home none`
13. edit your `/etc/fstab`, removing the old entry for `/home` and adding `/dev/mapper/home /home ext3 defaults 1 2`
14. verify your `fstab` entry: `mount /home`
15. restore default SELinux security contexts: `/sbin/restorecon -v -R /home`
16. reboot: `shutdown -r now`
17. The entry in `/etc/crypttab` makes your computer ask your luks passphrase on boot
18. Login as root and restore your backup

3.7.4. What you have just accomplished.

Congratulations, you now have an encrypted partition for all of your data to safely rest while the computer is off.

3.7.5. Links of Interest

For additional information on LUKS or encrypting hard drives under Fedora please visit one of the following links:

- [LUKS - Linux Unified Key Setup](#)⁵
- [HOWTO: Creating an encrypted Physical Volume \(PV\) using a second hard drive, pvmove, and a Fedora LiveCD](#)⁶

3.8. 7-Zip Encrypted Archives

[7-Zip](#)⁷ is a cross-platform, next generation, file compression tool that can also use strong encryption (AES-256) to protect the contents of the archive. This is extremely useful when you need to move data between multiple computers that use varying operating systems (i.e. Linux at home, Windows at work) and you want a portable encryption solution.

3.8.1. 7-Zip Installation in Fedora

7-Zip is not a base package in Fedora, but it is available in the software repository. Once installed, the package will update alongside the rest of the software on the computer with no special attention necessary.

3.8.2. Step-by-Step Installation Instructions

- Open a Terminal: Click `'Applications'` -> `'System Tools'` -> `'Terminal'`

⁷ <http://www.7-zip.org/>

- Install 7-Zip with sudo access: `sudo yum install p7zip`
- Close the Terminal: `exit`

3.8.3. Step-by-Step Usage Instructions

By following these instructions you are going to compress and encrypt your "Documents" directory. Your original "Documents" directory will remain unaltered. This technique can be applied to any directory or file you have access to on the filesystem.

- Open a Terminal: Click `'Applications'` -> `'System Tools'` -> `'Terminal'`
- Compress and Encrypt: (enter a password when prompted) `7za a -mhe=on -ms=on -p Documents.7z Documents/`

The "Documents" directory is now compressed and encrypted. The following instructions will move the encrypted archive somewhere new and then extract it.

- Create a new directory: `mkdir newplace`
- Move the encrypted file: `mv Documents.7z newplace`
- Go to the new directory: `cd newplace`
- Extract the file: (enter the password when prompted) `7za x Documents.7z`

The archive is now extracted into the new location. The following instructions will clean up all the prior steps and restore your computer to its previous state.

- Go up a directory: `cd ..`
- Delete the test archive and test extraction: `rm -r newplace`
- Close the Terminal: `exit`

3.8.4. Things of note

7-Zip is not shipped by default with Microsoft Windows or Mac OS X. If you need to use your 7-Zip files on those platforms you will need to install the appropriate version of 7-Zip on those computers. See the 7-Zip [download page](#)⁸.

GNOME's File Roller application will recognize your .7z files and attempt to open them, but it will fail with the error `"An error occurred while loading the archive."` when it attempts to do so. This is because File Roller does not currently support the extraction of encrypted 7-Zip files. A bug report (http://bugzilla.gnome.org/show_bug.cgi?id=490732 Gnome Bug 490732]) has been submitted.

3.9. Using GNU Privacy Guard (GnuPG)

GPG is used to identify yourself and authenticate your communications, including those with people you don't know. GPG allows anyone reading a GPG-signed email to verify its authenticity. In other words, GPG allows someone to be reasonably certain that communications signed by you actually are from you. GPG is useful because it helps prevent third parties from altering code or intercepting conversations and altering the message.

⁸ <http://www.7-zip.org/download.html>

3.9.1. Creating GPG Keys in GNOME

Install the Seahorse utility, which makes GPG key management easier. From the main menu, select System > Administration > Add/Remove Software and wait for PackageKit to start. Enter Seahorse into the text box and select the Find. Select the checkbox next to the "seahorse" package and select "Apply" to add the software. You can also install Seahorse at the command line with the command `su -c "yum install seahorse"`.

To create a key, from the "Applications > Accessories" menu select "Passwords and Encryption Keys", which starts the application Seahorse. From the "Key" menu select "Create New Key..." then "PGP Key" then click "Continue". Type your full name, email address, and an optional comment describing who are you (e.g.: John C. Smith, jsmith@example.com, The Man). Click "Create". A dialog is displayed asking for a passphrase for the key. Choose a strong passphrase but also easy to remember. Click "OK" and the key is created.



Warning

If you forget your passphrase, the key cannot be used and any data encrypted using that key will be lost.

To find your GPG key ID, look in the "Key ID" column next to the newly created key. In most cases, if you are asked for the key ID, you should prepend "0x" to the key ID, as in "0x6789ABCD". You should make a backup of your private key and store it somewhere secure.

3.9.2. Creating GPG Keys in KDE

Start the KGpg program from the main menu by selecting Applications > Utilities > Encryption Tool. If you have never used KGpg before, the program walks you through the process of creating your own GPG keypair. A dialog box appears prompting you to create a new key pair. Enter your name, email address, and an optional comment. You can also choose an expiration time for your key, as well as the key strength (number of bits) and algorithms. The next dialog box prompts you for your passphrase. At this point, your key appears in the main KGpg window.



Warning

If you forget your passphrase, the key cannot be used and any data encrypted using that key will be lost.

To find your GPG key ID, look in the "Key ID" column next to the newly created key. In most cases, if you are asked for the key ID, you should prepend "0x" to the key ID, as in "0x6789ABCD". You should make a backup of your private key and store it somewhere secure.

3.9.3. Creating GPG Keys Using the Command Line

Use the following shell command: `gpg --gen-key`

This command generates a key pair that consists of a public and a private key. Other people use your public key to authenticate and/or decrypt your communications. Distribute your public key as widely as possible, especially to people who you know will want to receive authentic communications from you, such as a mailing list. The Fedora Documentation Project, for example, asks participants to include a GPG public key in their self-introduction.

A series of prompts directs you through the process. Press the Enter key to assign a default value if desired. The first prompt asks you to select what kind of key you prefer:

Please select what kind of key you want: (1) DSA and ElGamal (default) (2) DSA (sign only) (4) RSA (sign only) Your selection? In almost all cases, the default is the correct choice. A DSA/ElGamal key allows you not only to sign communications, but also to encrypt files.

Next, choose the key size: minimum keysize is 768 bits default keysize is 1024 bits highest suggested keysize is 2048 bits What keysize do you want? (1024) Again, the default is sufficient for almost all users, and represents an "extremely" strong level of security.

Next, choose when the key will expire. It is a good idea to choose an expiration date instead of using the default, which is "none." If, for example, the email address on the key becomes invalid, an expiration date will remind others to stop using that public key.

Please specify how long the key should be valid. 0 = key does not expire d = key expires in n days w = key expires in n weeks m = key expires in n months y = key expires in n years Key is valid for? (0)

Entering a value of 1y, for example, makes the key valid for one year. (You may change this expiration date after the key is generated, if you change your mind.)

Before the `gpgcode>` program asks for signature information, the following prompt appears: Is this correct (y/n)? Enter `ycode>` to finish the process.

Next, enter your name and email address. Remember this process is about authenticating you as a real individual. For this reason, include your real name. Do not use aliases or handles, since these disguise or obfuscate your identity.

Enter your real email address for your GPG key. If you choose a bogus email address, it will be more difficult for others to find your public key. This makes authenticating your communications difficult. If you are using this GPG key for `[[DocsProject/SelfIntroduction| self-introduction]]` on a mailing list, for example, enter the email address you use on that list.

Use the comment field to include aliases or other information. (Some people use different keys for different purposes and identify each key with a comment, such as "Office" or "Open Source Projects.")

At the confirmation prompt, enter the letter O to continue if all entries are correct, or use the other options to fix any problems. Finally, enter a passphrase for your secret key. The `gpg` program asks you to enter your passphrase twice to ensure you made no typing errors.

Finally, `gpg` generates random data to make your key as unique as possible. Move your mouse, type random keys, or perform other tasks on the system during this step to speed up the process. Once this step is finished, your keys are complete and ready to use:

```
pub 1024D/1B2AFA1C 2005-03-31 John Q. Doe (Fedora Docs Project) <jqdoe@example.com>
Key fingerprint = 117C FE83 22EA B843 3E86 6486 4320 545E 1B2A FA1C
sub 1024g/CEA4B22E 2005-03-31 [expires: 2006-03-31]
```

The key fingerprint is a shorthand "signature" for your key. It allows you to confirm to others that they have received your actual public key without any tampering. You do not need to write this fingerprint down. To display the fingerprint at any time, use this command, substituting your email address: `gpg --fingerprint jqdoe@example.com`

Your "GPG key ID" consists of 8 hex digits identifying the public key. In the example above, the GPG key ID is 1B2AFA1C. In most cases, if you are asked for the key ID, you should prepend "0x" to the key ID, as in "0x1B2AFA1C".



Warning

If you forget your passphrase, the key cannot be used and any data encrypted using that key will be lost.

3.9.4. About Public Key Encryption

1. [Wikipedia - Public Key Cryptography](#)⁹
2. [HowStuffWorks - Encryption](#)¹⁰

General Principles of Information Security

The following general principals provide an overview of good security practices:

- encrypt all data transmitted over networks to help prevent man-in-the-middle attacks and eavesdropping. It is important to encrypt authentication information, such as passwords.
- minimize the amount of software installed and running services.
- use security-enhancing software and tools, for example, Security-Enhanced Linux (SELinux) for Mandatory Access Control (MAC), Netfilter iptables for packet filtering (firewall), and the GNU Privacy Guard (GnuPG) for encrypting files.
- if possible, run each network service on a separate system to minimize the risk of one compromised service being used to compromise other services.
- maintain user accounts: create and enforce a strong password policy; delete unused user accounts.
- routinely review system and application logs. By default, security-relevant system logs are written to `/var/log/secure` and `/var/log/audit/audit.log`. Note: sending logs to a dedicated log server helps prevent attackers from easily modifying local logs to avoid detection.
- never log in as the root user unless absolutely necessary. It is recommended that administrators use **sudo** to execute commands as root when required. Users capable of running **sudo** are specified in `/etc/sudoers`. Use the **visudo** utility to edit `/etc/sudoers`.

4.1. Tips, Guides, and Tools

The United States' *National Security Agency (NSA)*¹ provides hardening guides and tips for many different operating systems, to help government agencies, businesses, and individuals secure their systems against attack. The following guides (in PDF format) provide guidance for Red Hat Enterprise Linux 5:

- [Hardening Tips for the Red Hat Enterprise Linux 5](#)²
- [Guide to the Secure Configuration of Red Hat Enterprise Linux 5](#)³

The *Defense Information Systems Agency (DISA)*⁴ provides documentation, checklists, and tests to help secure your system (*Information Assurance Support Environment*⁵). The *UNIX SECURITY TECHNICAL IMPLEMENTATION GUIDE*⁶ (PDF) is a very specific guide to UNIX security - an advanced knowledge of UNIX and Linux is recommended before reading this guide.

The DISA *UNIX Security Checklist Version 5, Release 1.16*⁷ provides a collection of documents and checklists, ranging from the correct ownerships and modes for system files, to patch control.

¹ <http://www.nsa.gov/>

⁴ <http://www.disa.mil/>

⁵ <http://iase.disa.mil/index2.html>

⁶ <http://iase.disa.mil/stigs/stig/unix-stig-v5r1.pdf>

⁷ http://iase.disa.mil/stigs/checklist/unix_checklist_v5r1-16_20090215.ZIP

Also, DISA has made available [UNIX SPR scripts](#)⁸ that allow administrators to check specific settings on systems. These scripts provide XML-formatted reports listing any known vulnerable settings.

⁸ <http://iase.disa.mil/stigs/SRR/unix.html>

Secure Installation

Security begins with the first time you put that CD or DVD into your disk drive to install Fedora. Configuring your system securely from the beginning makes it easier to implement additional security settings later.

5.1. Disk Partitions

The NSA recommends creating separate partitions for `/boot`, `/`, `/home`, `/tmp`, and `/var/tmp`. The reasons for each are different and we will address each partition.

`/boot` - This partition is the first partition that is read by the system during boot up. The boot loader and kernel images that are used to boot your system into Fedora are stored in this partition. This partition should not be encrypted. If this partition is included in `/` and that partition is encrypted or otherwise becomes unavailable then your system will not be able to boot.

`/home` - When user data (`/home`) is stored in `/` instead of in a separate partition, the partition can fill up causing the operating system to become unstable. Also, when upgrading your system to the next version of Fedora it is a lot easier when you can keep your data in the `/home` partition as it will not be overwritten during installation. If the root partition (`/`) becomes corrupt your data could be lost forever. By using a separate partition there is slightly more protection against data loss. You can also target this partition for frequent backups.

`/tmp` and `/var/tmp` - Both the `/tmp` and the `/var/tmp` directories are used to store data that doesn't need to be stored for a long period of time. However if a lot of data floods one of these directories it can consume all of your storage space. If this happens and these directories are stored within `/` then your system could become unstable and crash. For this reason, moving these directories into their own partitions is a good idea.

5.2. Utilize LUKS Partition Encryption

Since Fedora 9, implementation of *Linux Unified Key Setup-on-disk-format*¹(LUKS) encryption has become a lot easier. During the installation process an option to encrypt your partitions will be presented to the user. The user must supply a passphrase that will be the key to unlock the bulk encryption key that will be used to secure the partition's data.

¹ http://fedoraproject.org/wiki/Security_Guide/9/LUKSDiskEncryption

Software Maintenance

Software maintenance is extremely important to maintaining a secure system. It is vital to patch software as soon as it becomes available in order to prevent attackers from using known holes to infiltrate your system.

6.1. Install Minimal Software

It is best practice to install only the packages you will use because each piece of software on your computer could possibly contain a vulnerability. If you are installing from the DVD media take the opportunity to select exactly what packages you want to install during the installation. When you find you need another package, you can always add it to the system later.

6.2. Plan and Configure Security Updates

All software contains bugs. Often, these bugs can result in a vulnerability that can expose your system to malicious users. Unpatched systems are a common cause of computer intrusions. You should have a plan to install security patches in a timely manner to close those vulnerabilities so they can not be exploited.

For home users, security updates should be installed as soon as possible. Configuring automatic installation of security updates is one way to avoid having to remember, but does carry a slight risk that something can cause a conflict with your configuration or with other software on the system.

For business or advanced home users, security updates should be tested and schedule for installation. Additional controls will need to be used to protect the system during the time between the patch release and its installation on the system. These controls would depend on the exact vulnerability, but could include additional firewall rules, the use of external firewalls, or changes in software settings.

6.3. Adjusting Automatic Updates

Fedora is configured to apply all updates on a daily schedule. If you want to change the how your system installs updates you must do so via "Software Update Preferences". You can change the schedule, the type of updates to apply or to notify you of available updates.

In Gnome, you can find controls for your updates at: System -> Preferences -> Software Updates. In KDE it is located at: Applications -> Settings -> Software Updates.

6.4. Install Signed Packages from Well Known Repositories

Software packages are published through repositories. All well known repositories support package signing. Package signing uses public key technology to prove that the package that was published by the repository has not been changed since the signature was applied. This provides some protection against installing software that may have been maliciously altered after the package was created but before you downloaded it.

Using too many repositories, untrustworthy repositories, or repositories with unsigned packages has a higher risk of introducing malicious or vulnerable code into your system. Use caution when adding repositories to yum/software update.

References

The following references are pointers to additional information that is relevant to SELinux and Fedora but beyond the scope of this guide. Note that due to the rapid development of SELinux, some of this material may only apply to specific releases of Fedora.

Books

SELinux by Example

Mayer, MacMillan, and Caplan

Prentice Hall, 2007

Tutorials and Help

Understanding and Customizing the Apache HTTP SELinux Policy

<http://fedora.redhat.com/docs/selinux-apache-fc3/>

Tutorials and talks from Russell Coker

<http://www.coker.com.au/selinux/talks/ibmtu-2004/>

Generic Writing SELinux policy HOWTO

<http://www.lurking-grue.org/writing/selinuxpolicyHOWTO.html>

Red Hat Knowledgebase

<http://kbase.redhat.com/>

General Information

NSA SELinux main website

<http://www.nsa.gov/selinux/>¹

NSA SELinux FAQ

<http://www.nsa.gov/selinux/info/faq.cfm>²

Fedora SELinux FAQ

<http://fedora.redhat.com/docs/selinux-faq-fc3/>

SELinux NSA's Open Source Security Enhanced Linux

<http://www.oreilly.com/catalog/selinux/>

Technology

An Overview of Object Classes and Permissions

http://www.tresys.com/selinux/obj_perms_help.html

Integrating Flexible Support for Security Policies into the Linux Operating System (a history of Flask implementation in Linux)

http://www.nsa.gov/research/_files/selinux/papers/selsymp2005.pdf

¹ <http://www.nsa.gov/research/selinux/index.shtml>

² <http://www.nsa.gov/research/selinux/faqs.shtml>

Chapter 7. References

Implementing SELinux as a Linux Security Module

http://www.nsa.gov/research/_files/publications/implementing_selinux.pdf

A Security Policy Configuration for the Security-Enhanced Linux

http://www.nsa.gov/research/_files/selinux/papers/policy/policy.shtml

Community

Fedora SELinux User Guide

<http://docs.fedoraproject.org/selinux-user-guide/>

Fedora SELinux Managing Confined Services Guide

<http://docs.fedoraproject.org/selinux-managing-confined-services-guide/>

SELinux community page

<http://selinux.sourceforge.net>

IRC

irc.freenode.net, #selinux, #fedora-selinux, #security

History

Quick history of Flask

<http://www.cs.utah.edu/flux/fluke/html/flask.html>

Full background on Fluke

<http://www.cs.utah.edu/flux/fluke/html/index.html>