



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Red Hat Linux 7.1 Installation Hardening Checklist

Paul Sery

The only way to reasonably secure your Linux workstation' is to use multiple layers of defense. There is no single system, such as a firewall or authentication process, that can adequately protect a computer. For instance, you may choose a good passwords and discover that an intruder has circumvented the login process altogether by exploiting a buffer overflow to gain root access. You must design a system consisting of multiple layers of defense.

This checklist is designed to help set up multiple layers of defense on a Red Hat Linux 7.1 workstation. The security measures described here are designed to work together to create a reasonable secure system. They work both internally and at the perimeter of a Red Hat Linux computer. When put together, the host computer is both reasonably safe and usable.

These measures are based on the author's experience from installing and running dozens of Red Hat Linux workstations. That experience includes reading and implementing vendor recommendations. Information and security insights have also been acquired from organizations such as SANS, USENIX, magazines such as Linux Journal and SysAdmin.

This checklist makes the following assumptions about the host machine:

- Desktop PC owned by the user's employer
- Primary purpose is to function as a workstation for one person (not a server)
- One or more disks with 2 GB or more of space
- Single Red Hat Linux 7.1 installation (no dual-boot)
- Workstation type installation (no individual package selections)
- Do not mount/automount remote NFS/Samba partitions
- Send/receive e-mail via client-server system based on imap/pop (no sendmail daemon)
- Use local router for NTP synchronization
- Access to local area network, with Internet connection

The measures described in this list will work on other configurations. However, it is optimized for the above workstation configuration. This workstation configuration is not as tight as you would use for a server configuration. A workstation must be easier to use than a server. Not as many people depend on a workstation as a server either so there is not as much at risk and security can be less tight.

Note: This checklist assumes that you are familiar with the Red Hat Linux installation process. It describes the overall configuration results but does not provide detailed instructions about what buttons to click and so on. If you need further installation assistance, please consult the Red Hat Installation guide at <http://www.redhat.com/support/manuals/RHL-7.1-Manual/install-guide>.

Please proceed to execute the following steps in order. For instance, if you connect to a live network before configuring and running Tripwire, you can not be 100% sure that the checksum database that it creates - and uses to check file integrity - will not be compromised.

1. Configure the BIOS
2. Install the Red Hat Linux 7.1 Workstation class
3. Remove unnecessary packages
4. Install Additional Packages
5. Update Red Hat packages
6. Stop Unnecessary Daemons
7. Secure the Linux File systems
8. Install a legal banner
9. Configure OpenSSH
10. Configure NTP
11. Configure or eliminate xinetd
12. Configure a host-based firewall based on iptables
13. Remove unnecessary login accounts
14. Monitor your log files
15. Modify networking parameters
16. Configure Tripwire
17. Connect to live network
18. Check your physical Security
19. Test your workstation
20. Further reading

Configure the BIOS

The rule-of-thumb is that anyone who has physical access to a computer can gain root access to that computer. This is especially true of the PC architecture because Linux, or other operating systems can be booted from a floppy or CD-ROM drive. However, you can prevent the casual intruder from booting your machine by setting a BIOS password.

For instance, on a DELL Precision 410, press the *F2* key while the system is powering up. From the BIOS menu, select the *Security* menu. Select the *Set Supervisor Password* menu and then enter a password. Next, turn on the *Password on boot:* option. Save the settings and continue booting the machine.

Please take note that if you set a BIOS password your machine will not be able to reboot if left unattended. The BIOS password must be entered from the machine's keyboard and, therefore you must be present.

The process for accessing the BIOS configuration menus can vary from PC to PC. The F1, F2, F12 and Del keys are typically used to access the BIOS menu.

Physically locking your workstation in an office or other enclosed area also adds to its security. Intruders can not gain access and, thus can not use the above mentioned tactics

to break in.

Install the Red Hat Linux 7.1 Workstation class

This document does not describe the Red Hat Linux installation process in detail (please consult Red Hat's documentation for step-by-step installation instructions). We describe the general parameters that should be used during the installation however. For instance, we choose to use the Workstation installation type, which affects many of the configuration choices described later.

The following list describes the choices we suggest you use for installing Red Hat Linux. The items follow the general flow of installation process.

- * Workstation installation type
- * Manual partition
- * Add the following partitions
 - / (root)
 - /usr
 - /var
 - /home

* You can vary partitions sizes depending on your needs. Their sizes, as well as the actual partition selection, are very much a matter of personal preference. However, there are certain limitations that should be recognized: the root (/) partition must be large enough to store all the kernel files that you expect to use. Otherwise, it should be as small as possible. The caveat is that the system stores temporary files to the /tmp directory which is part of the root file system. One solution is to create a link from /tmp to /var/tmp so that the temporary files are stored in the /var partition. (This solution occasionally causes problems if the /var partition can not be mounted at boot time for any reason. In that case, you need to boot into single user mode and manually run fsck on /var.) The other solution is to create a separate /tmp filesystem.

The /usr partition stores most of the application and utilities. It must be large enough to store all of system application, library and various system files.

The /var partition must be large enough to store all of the variable information that you expect to generate on your system. Those files include all of the spool and log files; it must also store the RPM package logs.

* We prefer to create one large storage area for optional application and data files rather than several smaller ones. For instance, traditional technique calls for creating separate /opt, /usr/local and /home file systems. We, however, combine those file systems into a single /home partition. We create soft links - /opt and /usr/local - that point to directories - /home/opt and /home/local - on the /home partition. The end result is more efficient use of available disk space without sacrificing any functionality.

Note that if you choose to consolidate the traditional /opt, /usr/local and /home partitions into a single one, then the name you choose for that single partition is arbitrary. We choose to use the name of /home because it is one of the standard names that the Red Hat Linux installation system provides, but you can name it anything you want. For instance, the partition name could be /local, /space or any one that you prefer.

* It is difficult to reasonably create more than two partitions on disks of 2 GB or smaller. It is better to allow the Red Hat installation system to choose the partitions. The installation system will create the following three partitions on a 2 GB disk:

- * swap 2.5 time RAM up to 256 MB
- * root (/) partition should be roughly 50 to 250 MB
(in general the smaller root is the better, which suggests the 50 MB size. However, the superuser home directory is /root; don't confuse the directory name with the root file system - they are two different entities. If you perform work as the superuser and store files in the /root directory, then you may want to allocate the larger size the root partition)
- * /home Remaining space

* A 4 GB disk permit you to create more partitions (use manual/Disk Druid):

- * swap 2.5 time RAM up to 256 MB
- * root (/) partition should be roughly 250 MB.
- * /usr 1.5 GB for average
- * /var 500 MB
- * /home Remaining space

* The following sizes are adequate for larger disks. Note that the root remains the same size as previous configuration and /usr expands by only 25%. There is generally little need to make those two partitions much larger than specified here unless you need to install very large, or very many, additional packages. Third party packages are installed into /usr/local and /opt which are linked to /home/local.

- * swap 2.5 time RAM up to 256 MB
- * root (/) partition should be roughly 250 MB.
- * /usr 2 GB for average
- * /var 1 GB
- * /home Remaining space

* Select your network parameters as dictated by your LAN configuration. Consult your local systems administrator for your IP address, netmask, DNS server and default gateway addresses.

* Select the *No firewall* option (we will be configuring a better, stateful firewall based on iptables later in this guide).

* Enter "good" root and user passwords. Good passwords should not contain words or phrases that can be found in any dictionary. You should include non-alphanumeric characters in the mix.

* You can select individual packages as desired. However, we let the installation

system select the individual packages and then remove certain ones later.

Remove unnecessary packages

When using the Red Hat Linux Workstation installation class, many packages are installed that are unnecessary and/or can cause security problems. You should examine the total installed package list in order to determine other packages to remove. Every package that you remove eliminates potential vulnerabilities. Of course, removing packages can also remove needed functionality so use caution.

To examine the packages installed on your workstation, use the command **rpm -qa** to list all of the packages. The **rpm -qi *package*** command provides information about individual packages. For instance, **rpm -qi iptables** displays information about the iptables package.

The following lists describes some of the packages that typically can be removed from a computer functioning as a personal workstation. For instance, the m4 package is used to create sendmail.cf configuration files. We have made the assumption that this machine uses a client/server e-mail client like Netscape Communicator and, thus does not run sendmail daemon. (sendmail can still be useful for sendmail out system generated e-mail, however, so we don't remove the sendmail package.) However, your individual configuration may require the use of M4 so use your best judgment.

Sample suggestions for Red Hat package removal:

- * m4 We don't expect to create sendmail.cf files
- * dhcpd We assume static IP addresses
- * ipchains We use the newer iptables/Netfilter system instead of ipchains
- * kernel-source We run standard linux kernel - do not need source

This list was been compiled in a very conservative manner and is just a suggestion. We have to be very careful when suggesting what packages can be removed because individual users have such varied needs. Use your best judgment when determining what packages to remove.

Install Additional Red Hat Packages

Install additional packages as required. For instance, you may want to install Sun Micro system's StarOffice Desktop productivity suite package. Download and install the packages from their providers as desired.

Update Red Hat packages

Red Hat updates their packages whenever security vulnerabilities are identified and corrected. It is essential to update the packages installed on your Linux computer as frequently as possible.

Red Hat provides an automatic update facility called up2date. If you subscribe to their Software Manager service, you can have your packages updated automatically. Information on the service is found at:

*ccu<http://www.redhat.com/support/manuals/RHNetwork/ref-guide/up2date.html>

Otherwise, proceed as follows:

1. Copy the updated packages from: [ftp.redhat.com/pub/redhat/redhat-7.1-en/os/RedHat/RPMS](ftp://ftp.redhat.com/pub/redhat/redhat-7.1-en/os/RedHat/RPMS)
and/or
<http://www.redhat.com/support/errata/rh71-errata-security.html>
into the `/usr/local/src` directory.
2. Run the command **`rpm -Uvh /usr/local/src/*`**

Periodically check the Red Hat web site for further updates.

Stop Unnecessary Daemons

Deleting the packages described in the section "Removing unnecessary packages" will prevent their daemons from starting. However, you may decide not to remove some or all of the packages in which case you should prevent the following daemons from running:

ipchains	default configuration uses iptables
portmap	default configuration uses no NFS and Samba file systems
nfslock	default configuration uses no ifs and Samba file systems
netfs	default configuration uses no NFS and Samba file systems
autofs	default configuration uses no automounting
apmd	power management generally only necessary on a laptop
isdn	default configuration uses local network connection (LAN)
pppoe	default configuration uses local network connection (LAN)
sendmail	default configuration uses client-server e-mail retrieval (imap/PPP)
gpm	we use X Window and this is not necessary
anacron	default configuration leaves power on

The scripts that start Red Hat Linux services are stored in the `/etc/rc.d/init.d` directory. However, the Red Hat boot process executes soft links that point to the scripts to start the services. The links are stored in directories corresponding to the system runlevel: `/etc/rc.d/rc3.d` is used to start the non-graphical services and `/etc/rc.d/rc5.d` for the graphical ones. Soft links that start with a capital "S" are used to start services. Links starting with a "K" kill, or stop, a server.

For instance, the soft link `/etc/rc.d/rc5.d/S08portmap` points to `/etc/rc.d/init.d/portmap` script and starts the portmapper when Red Hat Linux boots into graphical - runlevel 5 - mode. Therefore, if you rename the file, then the portmap script will never be executed.

Rename the corresponding links as follows:

1. `mv S08ipchains .Nostart.S08ipchains`
2. `mv S13portmap .Nostart.S13portmap`
3. `mv S14nfslock.Nostart.S14nfslock`

4. ___ mv S25netfs .Nostart.S25netfs
5. ___ mv S26apmd .Nostart.S26apmd (skip if using a laptop)
6. ___ mv S80isdn .Nostart.S80idsn
7. ___ mv S80pppoe .Nostart.S80pppoe
8. ___ mv S80sendmail .Nostart.S80sendmail (skip if no imap/PPP)
9. ___ mv S85gpm .Nostart.S85gpm
10. ___ mv S95anacron .Nostart.S95anacron (skip if using a laptop)

If you need to run any of these services in the future, rename to their original state.

Alternative methods controlling system services is available via the `chkconfig` and `ntsysv` utilities. For instance, run the command **`chkconfig --list`** to see all the possible services and their state. To turn off a service at all run-levels, run the command **`chkconfig <service name> off`**; turn the service on by running **`chkconfig <service name> on`**. The utility can also control services at individual run-levels by using the `--level` option. For example, to turn off the portmapper at run-level 5, run the command **`chkconfig --level 5 portmap off`**; to turn off the service at all run-levels use the command **`chkconfig portmap off`**.

The `ntsysv` utility provides the same functionality as `chkconfig`. Unlike `chkconfig`, it provides a simple graphical interface.

We prefer to manually rename the soft link files. The primary advantage is that a record of the original soft link is maintained. The original file name can be restored and the service will be automatically started at boot again. (The service will always be available because neither method touches the script stored in `/etc/rc.d/init.d`. For instance manually renaming `S13portmap` or using **`chkconfig portmap off`** leaves `/etc/rc.d/init.d/portmap` script in place. You can start or stop the script by running the command `/etc/rc.d/init.d/portmap start` or `/etc/rc.d/init.d/portmap stop`.) However, maintaining the soft link is not imperative and so use the `chkconfig` utility if you prefer.

Secure the Linux File Systems

Security is enhanced by making partitions non-writable and preventing set-uid programs/scripts from running. Non-writable partitions prevent intruders from introducing unauthorized files (a Trojan for instance) onto a computer. Disallowing set user id (setuid) prevents one user from executing a script or program as another user (if the owner is root, then the program is run as root).

Edit the `/etc/fstab` file as follows:

1. ___ Change the `/usr default` entry to `ro *`
2. ___ Change the `/var default` entry to `nosuid *`
3. ___ Change `/ (root) partition default` entry to `nosuid`
4. ___ Change `/home default` entry to `nosuid`
5. ___ Change the `default` entry to `nosuid` on other partitions like `/tmp`, `/opt`, etc as necessary.

* skip this step if you let the Red Hat Linux installation process automatically configure the file systems.

Please note that you increase the work require for adding, deleting and modifying RPM packages, and other software, by making the /usr partition read-only. Most RPM packages install files and directory to the /usr file system. You have to re-mount the /usr file system as read-write (rw) in order to add, delete and modify RPM packages. Remounting requires taking the system to single-user mode because many system processes access files on /usr; a file system can not be unmounted while processes are using it.

Remounting the /usr file system is a fairly difficult process. It may not seem to be worth the extra work. However, remember that it is much more difficult to re-constitute a compromised workstation. Making /usr read-only is a very effective security method and is worth the extra effort.

Install a legal banner

If you ever need to prosecute an intruder, it is necessary to warn him/her before a break-in occurs. It is necessary to install some form of legal banner on your computer. It is beyond our abilities as system managers and can not advise on the legal aspects of such a banner.

The following example legal notice is only suggestion. We do not suggest or recommend that you use it. We are not lawyers! Please consult a legal authority before composing your own message.

WARNING NOTICE TO USERS

This computer system is the property _____ . It is for authorized use only. Users (authorized or unauthorized) have no explicit or implicit expectation of privacy

Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized _____, and law enforcement personnel, as well as authorized officials of other agencies, both domestic and foreign. By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of _____ personnel

Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties

By logging on to this system, you indicate your awareness of and agreement with the terms and conditions of this warning notice. If you do not agree with these terms and conditions DISCONTINUE all efforts to access or

utilize _____ computing equipment and information

Intall the your banner as follows:

1. ___ Create a text file containing the legal banner. For instance, /usr/local/etc/banner.txt
2. ___ Make a backup of the /etc/issue and /etc/motd files.
`cp /etc/issue /etc/issue.orig`
`cp /etc/motd /etc/motd.orig`
3. ___ Copy the banner text into the /etc/issue and /etc/motd files.
`cp /usr/local/etc/banner.txt /etc/issue`
`cp /usr/local/etc/banner.txt /etc/motd`

If and when you Telnet to your workstation, the banner will be displayed before you enter your username and password. If you use SSH to connect, the banner will be displayed after you are authenticated.

4. ___ (Optional) If you run an FTP server, edit the /etc/ftppass file to contain the following line.

```
banner /etc/issue
```

The FTP server is directed to display the legal banner notice before each login.

5. ___ If you are using the Red Hat Linux graphical run-level @md 5 @md then you need to set up graphical login screen to show the banner. Edit the /etc/X11/xdm/Xsetup_0 configuration file to include the following code

```
if [ ! -f /usr/local/etc/Xlegal ]; then
    echo "ERROR: Legal Notice Banner missing from /usr/local/etc/Xlegal"
exit 1
fi
/usr/X11R6/bin/xmessage -file /usr/local/etc/Xlegal -buttons " \
consent:3,:4,reject:6" -default "reject" -center
```

This code directs the X server to prompt users to agrees or disagree to the legal notice when logging in from the console.

6. ___ Restart the X server by either rebooting, changing run-levels or pressing the Ctrl, Alt and Backspace keys at the same time. Please be aware, that all your current that are running under your X session will be killed.

Configure OpenSSH

The Red Hat Workstation installation installs OpenSSH by default. OpenSSH is an open source version of the commercial Secure Shell (SSH) package that encrypts network communication; it also can authenticate both the client and user by passing public keys between the client and server.

You should use OpenSSH for all of your interactive shell sessions and file transfers if possible. You can connect to any SSH server with the OpenSSH client - SSH. To connect to your workstation, you must set up the OpenSSH server as follows:

1. ___ Generally, OpenSSH is installed by default during the Red Hat installation process. However, if it was not installed, then do so now. (You can check the installation status by running the command **rpm -qa | grep -i openssh**). If you do not see the OpenSSH-server package displayed, then you need to install it.)

```
rpm -ivh /mnt/cdrom/RedHat/RPMS/openssh-server*
```

2. ___ Modify the `/etc/ssh/sshd_config` file parameters shown below (you can leave the remaining parameters unchanged):

3. ___ Uncomment the following line. It makes the SSH server listed to all interfaces for incoming connection.

```
ListenAddress 0.0.0.0
```

4. ___ Deny root logins. Direct root logins should never be allowed. Force users to login as themselves first, and then `su -` to root. This leaves a better audit trail and also requires people to know two passwords instead of one.

```
PermitRootLogin no
```

5. ___ You can optionally force all logins to use the public/private key system. This authenticates the calling client as well as the user. Otherwise, user passwords are encrypted and used to authenticate all new connections but the host that the user is connecting from is never authenticated. That goes the same for the server that the user is connecting to - in this case the new Red Hat Linux box - is not authenticated to the client. This makes it possible for a man-in-the-middle attack to take place.

(Optionally, create a GNOME or KDE interface for your passkey.)

```
PasswordAuthentication no
```

6. ___ Once you have finished configuring the `/etc/ssh/sshd_config` file, restart the OpenSSH

```
/etc/rc.d/init.d/sshd restart
```

Please consult the following URLs for more information on the theory and operation of OpenSSH:

*[ccuhttp://www.redhat.com/support/manuals/RHL-7.1-Manual/customization-guide/openssh-clients.html](http://www.redhat.com/support/manuals/RHL-7.1-Manual/customization-guide/openssh-clients.html)

*[ccuhttp://www.redhat.com/support/manuals/RHL-7.1-Manual/customization-guide/openssh-servers.html](http://www.redhat.com/support/manuals/RHL-7.1-Manual/customization-guide/openssh-servers.html)

*[ccuhttp://www.redhat.com/support/manuals/RHL-7.1-Manual/ref-guide/s1-ssh-configfiles.html](http://www.redhat.com/support/manuals/RHL-7.1-Manual/ref-guide/s1-ssh-configfiles.html)

Configure NTP

We assume that you have access to a router that speaks the NTP protocol. It is a simple process to use the ntpd update your Linux computer's system time by creating a cron job to periodically query the router. The following instructions describe the configuration process.

1. ___ Log in as root
2. ___ Install the NTP client software: `rpm -ivh /mnt/cdrom/RedHat/RPMS/ntp*`
3. ___ Modify the crontab file to update the clock four times per day:
4. ___ Add the line `0 0,6,12,18 * * * /sbin/ntpdate -s {router ip address}`
5. ___ Save the changes. Cron will automatically run ntpdate four times per day.

Alternatively, you can run the daemon:

Under this configuration, ntpdate will run 4 times per day at midnight, 6 AM, noon and 6 PM. You may need to run ntpdate more frequently if your system time can not be kept within one second of the router. If you ever need to go to court to prosecute an intruder, all computers within a LAN should be within a second of each other. If your systems are not coordinated, then it is possible that the validity of your log records could be questioned. There are other situations where one-second resolution can be important such as when several people are collaborating together.

Using the ntpdate program is very simple. We recommend its use because it is adequate for most purposes and does not require running the ntpd daemon, which is more difficult to configure and maintain.

(If you want to run the ntpd daemon and have three local time sources, then please consult the NTP HOWTO available in the /usr/share/doc/ntp* directory. The following URLs provide additional information:

*ccuhttp://www.eecis.udel.edu/~ntp/ntp_spool/html/ntpd.htm

*ccu<http://www.europe.redhat.com/documentation/mini-HOWTO/Clock-3.php3#ss3.3>

Configure or eliminate xinetd (optional)

Red Hat Linux 7.0 introduced the xinetd system. The xinetd system replaces inetd. Xinetd provides its own access control and does not require tcp_wrappers. It also provides logging capabilities.

Red Hat installs xinetd by default. However, this configuration check list does not use any of the services that xinetd controls; we use OpenSSH to provide the interactive functionality that xinetd is designed to provide - for instance Telnet or FTP. Remove xinetd unless you configure your workstation to provide services like Telnet. Remove the xinetd RPM package as follows:

```
rpm -e xinetd
```

If you do use xinetd controlled services, use the following instructions to tighten the xinetd configuration.

Xinetd uses a single configuration file found in /etc/xinetd.conf. Each internet service is controlled by an individual configuration file found in the /etc/xinetd.d directory.

The xinetd package configures the following services by default.

```
/etc/xinetd.d/chargen
/etc/xinetd.d/chargen-udp
/etc/xinetd.d/daytime
/etc/xinetd.d/daytime-udp
/etc/xinetd.d/echo
/etc/xinetd.d/echo-udp
/etc/xinetd.d/time
/etc/xinetd.d/time-udp
```

These services are rarely used and should, in general, not be provided to your network. Each service is by default turned off by the following line in its configuration file.

```
disable = yes
```

We find it marginally better to remove or rename the files because it is obvious which services are not offered if the file does not exist.

Access control can be set in the general /etc/xinetd.conf file or from the individual service files found in /etc/xinetd.d. The access options are:

```
* only_from
* no_access
```

These options take IP addresses and/or network names as their parameters. For instance, if you only want to allow access from a local network, and no where else, then use the following setup:

```
only_from = 192.168.1.0/24
```

(the "/24" notation specifies a class C address space; "/16" indicates a class B and "/8" a class A address space.)

Alternatively, you can use the no_access option to allow access from everywhere but the given locations. The following example allows everyone but the local network and a specific host to gain access:

```
no_access = 192.168.1.0/24 mybox.mynet.com
```

Use these *only_from* and *no_access* options in the */etc/xinetd.conf* file if you want the restrictions to work on all services. Otherwise, place the instructions in the particular service configuration file. For instance, if you are running a tenet server and want to allow connection from one subnet only, then place the following line in the */etc/xinetd.d/telnet* file:

```
only_from = 192.168.1.0/24
```

Don't forget to turn off the disable function.

```
disable = no
```

Restart the *xinetd* daemon.

```
/etc/rc.d/init.d/xinetd restart
```

Please consult the following URLs for more information on the theory and operation of *xinetd*:

*[ccuhttp://www.macsecurity.org/resources/xinetd/tutorial.shtml](http://www.macsecurity.org/resources/xinetd/tutorial.shtml)

*[ccuhttp://www.xinetd.org](http://www.xinetd.org)

*[ccuhttp://www.europe.redhat.com/documentation/rhl7/ref-guide-en/s1-sysadmin-access.php3](http://www.europe.redhat.com/documentation/rhl7/ref-guide-en/s1-sysadmin-access.php3)

Configure your host-based firewall using iptables

We assume that your private network is protected by a firewall at the Internet gateway. Such a firewall will protect you from many of the dangers external to your network. However, you can increase your protection significantly if you run a firewall at your machine's network interface.

A host-based firewall helps protect you against the insider threat and also misconfigured or inadequate general-purpose firewalls. There are two types of firewalls: IP filtering and proxy based. IP filters are generally the easier of the two to configure. They provide very good protection when allowing only outgoing connections; they generally provide adequate protection when providing mostly outgoing connectivity. For these reasons we recommend using an IP filtering firewall.

IP filters examine the source and destination addresses and port of every IP packet that passes through your network interface(2). IP filters accept or deny packets based on a set of rules that you design.

Iptables is a state-full IP filtering system. (Note that the system consists of the parts: Netfilter works at the kernel level and performs that actions of the IP packets and iptables is the user space system used to control Netfilter.) It is capable of making decisions based on both the origin/destination of a packet and its state. A packet's state is dependent on

whether it is part of an existing (or new) connection.

The following rule-set describes how to configure a host-based firewall based on iptables. This firewall uses the philosophy of denying all packets to begin with. Individual openings are created to allow certain types of connections. For instance, we create an opening to allow incoming Secure Shell connections.

```
IPTABLES="/sbin/iptables"

# get rid of any existing chains
$IPTABLES --flush
$IPTABLES --flush -t nat

# deny all traffic to start
$IPTABLES --policy INPUT    DROP
$IPTABLES --policy OUTPUT   DROP
$IPTABLES --policy FORWARD  DROP

# allow all internal traffic
$IPTABLES -A OUTPUT -j ACCEPT -o lo
$IPTABLES -A INPUT  -j ACCEPT -i lo

# allow all outgoing TCP/UDP connections
$IPTABLES -A OUTPUT -p tcp -m state --state NEW,ESTABLISHED,RELATED \ -j ACCEPT
$IPTABLES -A OUTPUT -p udp -m state --state NEW,ESTABLISHED,RELATED \ -j ACCEPT
$IPTABLES -A INPUT  -p tcp -m state --state ESTABLISHED,RELATED -j \ ACCEPT
$IPTABLES -A INPUT  -p udp -m state --state ESTABLISHED,RELATED -j \ ACCEPT

# allow incoming SSH connections
$IPTABLES -A INPUT -p tcp -m state --state NEW,ESTABLISHED,RELATED \ --dport 22 -j ACCEPT
```

These rules create a firewall that functions essentially as a diode. Most connections can be made from your machine to the outside world, but only Secure Shell connections can be made from the outside in. Your machine's security is greatly enhanced. Many existing vulnerabilities, for instance a theoretical buffer overflow in your print daemon, can not be taken advantage of. Packets can never be delivered to the vulnerable service to exploit them.

Add any or all of the following rules to use additional services.

```
# allow X connections
$IPTABLES -A INPUT -j ACCEPT -i eth0 -p tcp --dport 6000:6004 -m state \ --state NEW,ESTABLISHED
```

Please consult the following URLs for more information on the theory and operation of iptables:

- * <http://netfilter.samba.org/unreliable-guides/>
- * http://www.linuxnewbie.org/nhf/intel/security/iptables_basics.html

Remove unnecessary login accounts

Several unnecessary login accounts installed by default. The following list provides suggestions about what accounts are unnecessary given the workstation configuration used in this paper. As usual, you may need to adjust this list depending on your particular configuration.

mail
news
uucp
games
gopher
ftp
nscd
mail
mailnull
ident
rpc
rpcuser
xfs
gdm

You can determine if any files/directories belong to a group by using the **find / -user *name*** command where name is the user name to be queried. For instance, running the command **find / -user ftp** locates all files that belong to the ftp user. You can investigate each user name in the /etc/passwd file to find every one that is not being used.

Use the **userdel** command to remove the accounts. For instance, to remove the rpc user, run the command: **userdel rpc**.

You can minimize the /etc/group entries too. Some suggested groups to remove are:

ftp
news
uucp
gopher
dip
nscd
gopher
dip
games
rpc
rpcuser

Use the command **find / -group *name***, where name is the group in question to verify that a group is not being used. For instance, **find / -group dip** locates all files and directories belonging to the dip group.

Use the `groupdel` to remove group names. For instance, the following example removes the `uucp` group: **`groupdel uucp`**.

Note that some files and directories will be owned by root but belong to their own group. For instance, the `/var/spool/mail` directory has this arraignment.

Monitor your log files

The Linux kernel, daemons and many applications generate messages that the `syslogd` daemon can record for later examination. Red Hat Linux, by default, logs much of the information generated by daemons. The default configuration should be adequate for a typical workstation.

It is important to monitor your logs. However, it is beyond the scope of this document to describe in detail what one looks for in log files. We can state that the more you look at your logs the more familiar you'll become with how your system operates and what patterns are normal. When you become familiar with your workstation's operational patterns you will be able to better recognize abnormalities. Abnormal operations can indicate security related problems.

The `/etc/syslog.conf` file controls the operation of the `syslogd` daemon. The `syslog.conf` file contains entries that take are divided into a selector and action. The selector takes the form of comma separated fields of "facility.level". The action parameter describes where the messages that match the selector are sent.

The "facility.level" format of the selector field describes the type of message and its priority to be logged. The facility field determines how the following types of processes are handled:

*	user	User space processes
*	kern	Kernel space processes
*	mail	E-mail based processes
*	daemon	System daemons
*	auth	Authentication based processes
*	lpr	Printer related processes
*	news	News related processes
*	uucp	UUCP based mail
*	cron	Messages generated by the cron system
*	local0-local7	Reserved for user-generated processes
*	*	The asterix is a wild card for all facilities

Each facility is organized by their level of importance. The log levels are:

*	emerg	Emergency messages that are generated by panic conditions and should be made known to all users. (emerg replaces the former panic level.)
---	-------	---

- * alert Urgent situations that need to be immediately corrected
- * crit Critical conditions such as hardware problems
- * err General errors
- * warning Warning messages
- * info Informational messages
- * debug Messages generated to assist in debugging
- * none Prevents facilities from generating messages.

Red Hat configures the syslogd daemon to save log information in the /var/log directory. Some of the important files are:

- * /etc/syslog.conf This file controls the syslogd daemon. It configures what type of information is saved to what log files.
- * /var/log/messages This file serves as a generic repository for log messages. By default, syslogd saves information of
- * /var/log/cron Messages generated by the cron daemon.
- * /var/log/secure Messages generated by failed login attempts are stored in this file.

The log files are rotated based on their size and age. Rotation is controlled by the /etc/logrotate.conf configuration file. It's contents are shown below.

```
# see "man logrotate" for details
# rotate log files weekly
weekly

# keep 4 weeks worth of backlogs
rotate 4

# send errors to root
errors root

# create new (empty) log files after rotating old ones
create

# uncomment this if you want your log files compressed
#compress

# RPM packages drop log rotation information into this directory
include /etc/logrotate.d

# no packages own lastlog or wtmp -- we'll rotate them here
/var/log/wtmp {
    monthly
    create 0664 root utmp
    rotate 1
```

}

Notice that log files are rotated on a weekly basis and four iterations are saved. Errors are e-mailed to the root user. When the current log file is rotated, it is renamed with a ".1" suffix and a new, empty file is created. For instance, current log entries are written to `/var/log/messages`. When that time arrives for rotating that file, it is renamed to `/var/log/messages.1` and a new `/var/log/messages` file is created. Previously rotated files are renamed from `messages.1` to `messages.2`, and the current `messages.4` is erased.

You can modify this sequence as desired. For instance, if you want log files rotated on a daily basis, change the weekly option to daily; use the monthly option to do the rotation every month.

Please consult the following sources for additional information on modifying and monitoring log files.

Modify Networking Parameters

There are several network configuration modifications that can be increase your workstations security:

- * Turn off IP forwarding
- * Disallow source routing
- * Make sure that routed is not installed

Use the following list to accomplish these tasks.

1. ___ Verify that `/etc/sysctl.conf` has the following line that turns off IP forwarding
`net.ipv4.ip_forward = 0`
2. ___ Verify that `/etc/sysctl.conf` has the following line that enables source route verification
`net.ipv4.conf.all.rp_filter = 1`
3. ___ Run the `sysctl` program to put those options in effect.
`sysctl -p /etc/sysctl.conf`
4. ___ Determine if the `routed` package is installed
`rpm -qa | grep -i routed`
5. ___ Remove `routed` if it is installed
`rpm -e routed`

Configure Tripwire

It is essential to conduct this step after completing the previous steps before connecting to a live network (the next step). You must create a reliable Tripwire database before connecting to your network. Once you connect to the network the odds go up that your machine could be compromised along with the Tripwire database. Having a Tripwire database created before connecting to a network ensures that it will not be corrupted by a network-based intrusion. Conducting this step after finishing the others, also simplifies

the task of creating the database. You'll have to re-initialize the database more than once if you create it earlier in the process.

Tripwire calculates MD5 check sums of specified files and directories. The checksums are used to detect changes made to your file system. Unauthorized changes can mean that an intrusion has occurred. Tripwire is a simple, but effective intrusion detection system (IDS).

Red Hat Linux installation system installs Tripwire by default. However, it uses a configuration and policy file that does not work well only when you install every package. Otherwise, Tripwire returns too many errors (because it is looking for files that have not been installed) to be useful. The following Tripwire configuration and policy files have been designed to work with the installation used here. Please view these files as a starting point because you will probably want to further customize them to your own system:

1. ___ Log in as root
2. ___ Run the installation script
 /etc/tripwire/twinstall.sh
3. ___ Enter the site keyfile pass phrase when requested. This pass phrase is used to encrypt the /etc/tripwire/site.key file.
4. ___ Enter the local keyfile pass phrase. The script generates a file whose name is your fully qualified domain name with the string "-local.key" appended to it. For instance, if your machine is me.mydomain.com, then the file will be named "me.mydomain.com-local.key".
5. ___ Enter the site key pass phrase from step 2. Tripwire authenticates you and then encrypts twcfg.txt file. The encrypted file is named tw.cfg.
6. ___ Enter the site key pass phrase from step 2. Tripwire authenticates you and then encrypts twpol.txt file. The encrypted file is named tw.pol.
7. ___ Edit the twpol.txt to fit your configuration. A sample policy file is shown below. It is straightforward and checks several system directories that should not change unless the system administrator adds, edits or deletes RPM packages or otherwise modifies the system files.

```
(
  rulename = "Invariant Directories",
  severity = $(SIG_MED)
)
{
  / -> $(SEC_INVARIANT) (recurse = 0) ;
  /home -> $(SEC_INVARIANT) (recurse = 0) ;
  /etc -> $(SEC_INVARIANT) (recurse = 0) ;
```

```

}
(
  rule name = "File System and Disk Administration Programs",
  severity = $(SIG_HI)
)
{
  /sbin          -> $(SEC_CRIT);
  /bin           -> $(SEC_CRIT);
  /usr          -> $(SEC_BIN);
}
(
  rulename = "Temporary directories",
  recurse = false,
  severity = $(SIG_LOW)
)
{
  /usr/tmp       -> $(SEC_INVARIANT);
  /var/tmp       -> $(SEC_INVARIANT);
  /tmp           -> $(SEC_INVARIANT);
}

(
  rulename = "User binaries",
  severity = $(SIG_MED)
)
{
  /usr/local     -> $(SEC_BIN);
  /opt           -> $(SEC_BIN);
}

(
  rulename = "Security Control",
  severity = $(SIG_HI)
)
/etc            -> $(SEC_CRIT);

# Libraries
(
  rulename = "Libraries",
  severity = $(SIG_MED)
)
/usr/lib        -> $(SEC_BIN);
/lib           -> $(SEC_BIN);
(
  rulename = "Critical system boot files",
  severity = $(SIG_HI)
)

```

```

)
{
  /boot          -> $(SEC_CRIT) ;
  !/boot/System.map ;
  !/boot/module-info ;
}
# These files change every time the system boots ##
(
  rulename = "System boot changes",
  severity = $(SIG_HI)
)
{
  /dev          -> $(SEC_CONFIG) ;
  /var/lock/subsys -> $(SEC_CONFIG) ;
  /lib/modules  -> $(SEC_CONFIG) ;
}

```

8. ___ Re-encrypt the new twpol.txt file as follows: twadmin --create-polfile twpol.txt

9. ___ Delete the un-encrypted nd twcfg.txt files.

```
rm -f /etc/tripwire/*.txt
```

10. ___ Initialize the tripwire database

```
twadmin --init
```

11. ___ Enter your local (not site) pass phrase

12. ___ Tripwire uses the tw.pol policy file to create check sums of the specified files and directories on your system.

13. ___ Do a quick check of your new system. For instance, rename a file found in the tripwire database.

```
mv /sbin/tune2fs /sbin/_tune2fs
```

14. ___ Run Tripwire in check mode

```
tripwire --check
```

15. ___ Tripwire should catch the change and display a message like the following:

```
-----
# Section: Unix File System
-----
```

```
-----
Rule Name: File System and Disk Administration Programs (/sbin)
Severity Level: 100
```

Added:
"/sbin/_tune2fs"

Removed:
"/sbin/tune2fs"

Modified:
"/sbin"

=====
Error Report:
=====

16. ___ Don't forget to rename the test file and re-initialize the Tripwire database.

```
mv /sbin/_tune2fs /sbin/tune2fs
tripwire --init
```

17. ___ The Tripwire database will need to be updated every time you add, delete or modify an RPM package. That is because our policy file is designed to act on entire directories instead of individual files. We believe that the extra work of re-initializing the Tripwire database pays off in the long term. The alternative - actually the default Tripwire policy file configuration is much more labor intensive in the long run because hundreds of lines are required to deal with the individual files. The alternative is also more dangerous in our opinion because it is easier to miss checking important files because the configuration file is so large. However, use your best judgment to determine which method is best for you.

Note that Red Hat RPMs typically install files into directories such as /bin, /sbin, /lib, etc. When you install new packages the Tripwire database must be updated to reflect the changes or else it will produce many false positives.

18. ___ Run the following command if you modify the twpol.txt database.

```
twadmin --create-polfile twpol.txt
```

19. ___ Repeat steps 10 through 12 to recreate the Tripwire database.

19. ___ Installing Tripwire creates a cron job (see /etc/cron.daily/tripwire-check) that runs Tripwire in check mode on a daily basis. The root user receives daily Tripwire reports via e-mail.

Note: Traditionally it has been necessary to protect the Tripwire database by storing it on read only media. If the database is ever compromised, then Tripwire can be made to give erroneous results allowing a break-in to go undetected. However, the correct Tripwire version protects the database by encrypting it. Thus, the information stored in the

database can not be altered unless the encryption is broken - an unlikely occurrence. If the database file itself is modified, then the signature will not match and the compromise will be detected.

More information on configuring Tripwire can be found at:

Tripwire, Inc. provides an on-line policy generating applet at:

*ccu<http://tpt.tripwire.com/tprc/servlet/VersionChoose?platform-name=RedHat+Linux>

Connect to a live network

Once you have implemented all of the security configurations described here, you can connect to a live network. It is essential that you save this step for the end of the process because it is possible for an intruder to take advantage of the installation process to gain access to your machine. Obviously, it is easier to break into an un-hardened system than after it has been configured correctly.

You will want to connect your workstation to your private network. Once connected reboot your machine to make sure that all of the changes you have made take effect.

Check your physical security

Any computer that can physically be accessed by unauthorized personnel can be broken into. There are numerous ways to break into a machine that you have access to. For instance:

1. ___ Booting from a floppy disc or CD-ROM drive. PCs built since the late 90's are generally capable of being booted directly from CD-ROM. Once the machine has been booted, the root (/) file system can be mounted and the root password erased. The machine can then be rebooted and root access is achieved.

This vulnerability can be mitigated by setting a BIOS password. However, that measure can be defeated by the following two steps.

2. ___ The boot disk can be removed from the machine and attached to another one. The disk can be mounted and the root password erased or changed. The disk's data can also be copied, destroyed or modified.
3. ___ The BIOS chip can be changed to remove the password.
4. ___ The power can be repeatedly cycled until the operating system fails to a root shell. The root shell is a fail-safe measure to allow the system administrator to manually fsck the disk. However, the unauthorized person can use this fail-safe mode to break into the machine as well. This vulnerability can be fixed by using a logged root file system. A logged file system is much more difficult to make fail.
5. ___ You can improve your physical security by locking your office - or wherever your workstation is located - whenever you leave for more than a short time.

Test your system

Conduct the following tests once you connect your workstation to a live network.

1. ___ You should not be able to boot from floppy or CD-ROM without entering the correct BIOS password.
2. ___ Remove all unnecessary Red Hat packages.
3. ___ Update Red Hat packages that have been identified as having security vulnerabilities.
4. ___ Verify that no unnecessary daemons are running.
5. ___ You should not be able to write to the /usr partition and that suid files can not be executed on the / (root), /home and /var partitions.
6. ___ You should perform all post Installation network modifications.
7. ___ Your legal banner should be displayed when you are logging into your workstation.
8. ___ Verify that Secure Shell is operating correctly.
9. ___ Verify that the NTP time cron job updates your clock regularly.
10. ___ Verify that xinetd correctly uses all access lists that you have configured.
11. ___ Verify that your IP filtering firewall is running correctly.
12. ___ Verify that the system log daemon is running.
13. ___ Verify that Tripwire is running.
14. ___ Verify that you have configured your physical security appropriately.

Further reading

The following paper provides a good overview on the subject of hardening Linux. It also contains several good references to related articles.

*ccu<http://www.sans.org/infosecFAQ/linux/hardening.htm>

Summary

Steps 5, 6, 9 and 16 are the most important steps in this list:

- * Step 5 is important because you must plug known security holes before they can be used against you. Automated systems can find common security holes on your workstation very effectively. The Ramen and Lion worms were exploited in this way.
 - * Step 6 removes unnecessary services and is a very cost-effective step.
 - * Step 9 configures one of the most important communication tools available. OpenSSH provides encrypted communication channel for not only your interactive shell sessions, but also for transferring files and proxying services such as X. OpenSSH also has the ability to authenticate both the user and host. Authenticating users is an essential part of security, but verifying the host that the user is connecting from/to is also important to avoid \u201cMan in the middle\u201d attacks.
 - * Step 16 creates an effective and easy to manage intrusion detection system.
- Finally, you need to configure Tripwire so that you have a snapshot of your virgin file

systems; once your system becomes active on the network you can never completely trust another snap shot, so it's essential to make it now.

You have reasonably secure workstation now that you have completed these steps. Please note that this is a good starting point but not a solution. Maintaining good security is an ongoing process that requires constant education and vigilance. The process involves updating software as security vulnerabilities are identified and corrected. It also involves monitoring your system and Tripwire logs. The better you know your workstation, the better you can protect it.

Please sign this document below. It certifies that you have taken reasonable precautions to protect your workstation.

Printed Name: _____

Signature: _____ Date: _____

© SANS Institute 2000 - 2005, Auth.