



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.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Introduction to Host Based Cyber Defense

Roy Nielsen

Giac Certification v1.4b

© SANS Institute 2005, Author retains full rights.

Table of Contents

Definitions	ii
Abstract	iv
Minimizing System Services	1
Minimizing Windows Services.....	1
Minimizing Unix Services.....	2
Unix meta-servers, inetd & xinetd.....	2
Unix Startup Scripts.....	3
Control Access.....	5
Password Control.....	5
Axiom of Least	5
BIOS Password.....	5
Boot loader Password	5
OS username and password	6
File Access.....	10
Windows file permissions.....	10
Unix file permissions.....	12
Controlling network access.....	14
Firewall.....	14
Unix Tcpwrappers	21
Unix “R” services	21
Logging	22
*nix Logging	22
Windows logging.....	22
Firewall logging.....	25
Windows BlackIce logging.....	25
Linux iptables logging	27
Solaris ipfilter logging	28
Server Logging	29
Filesystem Integrity Logging.....	29
Tripwire	29
Scanning for Malware	31
Patching	33
Applications.....	33
Operating Systems.....	34
Windows.....	34
Linux.....	34
Solaris.....	36
Special cases.....	36
Data Security	38
Stegonography	38
Encryption	38
Resources.....	38
Backups	39
Software.....	39
Hardware	40
Disaster Recovery Plan.....	43
Keep Up With Computer Security Technology and News	44
Host Based Cyber Defense Strategy Review.....	45
Bibliography	46
Web sites	47

Definitions

*nix: Generic term for a Unix based system. Some examples are Linux, Solaris, *BSD, HP-UX, AIX, SCO and Irix.

*BSD: Unix versions based on Berkley Standard Unix, such as current free versions – NetBSD, OpenBSD, DragonFly BSD and FreeBSD.

ACL: Access Control List – a method or rule for an operating system to determine security privileges.

Adware: Software that is installed to track what the user/victim does on the world wide web with the intention of targeting advertisements to the user/victim.

Application: A computer program that provides a human interface for organizing and storing data, computing tasks, processing data. Popular applications include Microsoft Office, Internet Explorer, Mozilla, Open Office, Adobe Acrobat, a multitude of fun and educational games.

BIOS: An acronym for Basic Input Output System. This small piece of software resides on a chip on the motherboard of the computer. When the computer is powered up, this chip sends signals to the CPU, memory, other chips, drives and computer cards attached to the system to put them in an initial working state.

GUI: Graphical User Interface

HIDS: Host Intrusion Detection System, software that will detect a change in the integrity of a computer device. Tripwire and Firewall logging can do this.

HowTo: A white paper written on “how to” set up or configuring a computer. There is a well known site for Linux HowTos called The Linux Documentation Project. This website can be found at www.tldp.org.

IDS: Intrusion Detection System, meant to detect a network attack.

IIS: Microsoft Internet Information Server.

IPS: Intrusion Prevention System, meant to protect a network from attack.

Malware: Any software installed on the system that deliberately compromises it. This includes trojans, spyware and adware.

Multi-platform: Works on more than one, often many, operating systems.

OEM: Original Equipment Manufacturer. Some examples are Dell, Sony, Gateway, HP, IBM, etc. They often sell their computers with a pre-installed OS to home users; this is often Windows XP Home Edition.

Operating System: Provides an interface and environment for computer programs to run. Windows, Solaris, Linux, OS X, Irix, Hpux, Aix, Sco and FreeBSD are Computer Operating systems.

Packet: a unit of network traffic.

Posix: Portable Operating System Interface for uniX. A generic specification for operating systems which many *nix operating systems follow at varying levels. A more complete definition can be found at: wi-fiplanet.webopedia.com/TERM/P/POSIX.html

RPM: RedHat Package Manager – a program that manages installation of files that have been pre-packaged in rpm format. For usage of the rpm command, see: www.rpm.org/max-rpm/

Server: A computer program that provides information or services to other applications. Popular web servers are Apache, Tomcat and IIS. Some Popular database servers are Mysql, Oracle, MSSql.

Spyware: Software that is installed on the computer without the user/victim's knowledge, with the intent of recording what the user/victim is doing.

ssh: an acronym for Secure SHell. www.openssh.com/ for *nix variants, and www.f-secure.com/products/fssh/ for Windows and some Unix variants. Ssh is a secure means of connecting to a shell on a remote computer. It uses secure sockets libraries to encrypt the connection so nothing is visible in plain text that goes through the secure connection.

Trojan: A computer program disguised as a friendly program (game, e-card, picture, MS word macro, etc) that executes commands and programs intended to harm the user in some way.

White paper: A research paper; in context of this paper, a research paper on computer security.

© SANS Institute

Abstract

There is a lot of attention given in the computer security community to network security. Viruses, trojans, spyware and other malware come from the computer network. IT departments often concentrate on network firewalls, IDS and IPS systems to protect their network.

Although these are important parts of a computer security plan, paying attention to the individual computer is also important to the overall computer security strategy. Several things can be done to secure the computer such as control access to the computer, practice logging, scan for malware, practice patching, categorizing data, practice encryption and perform regular backups. Just as important as any of the above is to keep abreast of current and upcoming computer security technology and news. A whole research paper or project could easily be written about each of the subjects in this paper. This is an overview of each topic and will raise the readers' awareness about host based security concerns. The spirit of this document is as the following Chinese proverb: "Feed a man a fish, feed him for a meal. Teach a man to fish, feed him for a lifetime".

© SANS Institute

Minimizing System Services

In the Unix world, a “service” is a facility “to perform a specific type of network task”¹. This document goes further to define a service to include daemons, which are servers that wait for a connection from another piece of software. Some examples of operating system services are:

Windows

- error reporting service
- help and support
- network location awareness
- remote registry
- terminal services

***nix**

- rexed
- sendmail
- smartd
- snmpd
- tftp

Any program that has an open socket waiting for a connection may be vulnerable to exploit. Even programs designed with security in mind have had problems. If a service is not required, disable it. The fewer services installed and/or running on the system, the less likely it is that an attacker will be successful in compromising the system.

Today's operating systems offer many services. By default, many services are started and are not required, or even used. For instance, on a Solaris 9 machine ftp, telnet, rsh and other services are on by default. Many machines do not need these services running to perform their everyday functions. If remote access is needed, ssh can be used, and with it, tools for remote copying of files in a secure way.

There are ways to stop these services, and by doing so, deny attackers the best and easiest ways into your system. Here are instructions to do this for some popular operating systems.

Minimizing Windows Services

There are several Windows XP services that are not generally required and should be stopped or set to manual. The way these services connect, transfer data, or authenticate are not safe. This can be done by going to the control panel -> admin tools -> computer management -> services and applications, go to the standard tab and sort by status (started).

An in depth guide on managing system services for Windows XP can be found at www.theeldergeek.com/services_guide.htm

¹ Frisch, Aileen, Essential System Administration, Second Edition. O'Reilly, 1995: 592 – 595

The following services should be set to manual, unless there is a specific need for the service:

- 1) Distributed Link Tracking -> Manual
- 2) Error Reporting Service -> Manual
- 3) Help and Support -> Manual
- 4) Messenger -> Manual (Also in C:\Program Files\Messenger – rename directory to “Messenger.off”)
- 5) Secondary Log on -> Manual
- 6) Upload Manager- Manual
- 7) Web Client – Manual
- 8) Wireless Zero Configuration – Manual

The following services should be set to disabled, unless there is a specific need for them:

- 1) Network Location Awareness -> Disabled
- 2) Remote Registry -> Disabled
- 3) SSDP Discovery Service -> Disabled
- 4) System Restore -> Disable (also in Control Panel->System->SystemRestore, make sure it's disabled here as well, it may already be)
- 5) Terminal Services -> Disabled
- 6) Universal Plug and Play - Disabled

Minimizing Unix Services

Just as with Windows, by default, there are running services that are not used, that should be turned off. With *nix, two areas need to be checked. First the “meta-server” that controls many of the internet related services requires configuration to minimize vulnerabilities then needs to be restarted. Second area to check is the startup scripts that initialize services at boot time.

Unix meta-servers, inetd & xinetd

inetd

Solaris and Slackware use `/etc/inetd.conf` to control system services. Some of the services started by this meta-server were written many years ago before security was a major concern for computer systems. Very few people need the services that are made available by the meta-server. For example in Solaris, comment out the services in the `/etc/inetd.conf` file except the following, which is needed for cdrom and tape support:

```
100155/1      tli      rpc/ticotsord  wait    root    /usr/lib/smedia/rpc.smsserverd
```

If the machine is a server that will host a print server for other machines to connect to, it will need the following service running:

```
printer stream tcp nowait    root    /usr/lib/print/in.lpd    in.lpd
```


When the `inetd.conf` file has been modified to be secure, the `inetd` service will need to be restarted. To accomplish this, do the following:

```
kill -s HUP <pid_of_inetd_service>
```

This will restart the `inetd` service without changing its process id or running date and is a reliable way to get the `inetd` process to re-read the configuration file.

xinetd

Fedora and RedHat use a different scheme. The `xinetd` services are controlled by one file per service in the `/etc/xinetd.d` directory. Most services in this directory are stopped by default. If changes need to be made, the `xinetd` service will need to be restarted as well.

```
kill -s HUP <pid_of_xinetd_service>
```

On Unix variant systems, some services that are dangerously insecure are `ftp`, `telnet`, `rsh` and other `r-services`. The best strategy, as stated above, is to not turn on a service unless it is specifically needed and can be reliably secured.

Unix Startup Scripts

Unix systems can boot into several modes, or run levels. Run level 1 is generally known as “single user mode”. In this mode, only one user may log in, and that is only on the console. For run levels 2 through 5 there is no consensus between variations of Unix. For Solaris, run level 2 is a limited multi-user environment, and run level 3 has full multi-user support, with NFS². In most versions of Linux, for instance, run level 5 starts the graphical interface³. In Solaris, however, run level 3 starts the graphical interface. Different services and servers are started in different run levels.

In Linux, when the system is being initialized, it goes directly to the run level defined in `/etc/inittab`. In Solaris, the initialization process goes through each run level until the desired run level is reached.

To disable a service or server from running, find the startup script and disable it. The actual startup scripts can be found in the `/etc/init.d`. Each run level has its own directory, in which the scripts are links that point to the `/etc/init.d` files. For instance, run level 3 scripts can be found in `/etc/rc.d/rc3.d`. The files in this directory will either start with a `K` or an `S`. The initialization process looks for the files that start with `S`, and starts them up. The `K` scripts can be used to “kill” or stop a service on some `*nix` systems. To disable a script, rename a script from `SXX<service-name>` to something like `disabled.SXX<service-name>`. Do not rename the script from `SXX<service-name>` to `KXX<service-name>` because this could overwrite a specifically

² Sun Microsystems, System Administration Guide: Basic Administration.
docs.sun.com/db/doc/817-6958/6mmafc30c?a=view

³ Wirzenius, Lars; Oja, Joanna et al, The Linux System Administrator's Guide Version 0.8. The Linux Documentation Project, 2003: 9.3 at the following website: www.tldp.org/LDP/sag/html/x2140.html

written kill script necessary for regular system use. Also when executing the “ls” command, the renamed disabled.SXX<service-name> file will be very easily distinguished.

Some startup scripts to disable in Solaris are:

S50httpd
S70uucp
S76snmpd
S90samba

Some startup scripts in Fedora Core 2, like the following services, do not need to be run on most machines: Apache (httpd), smartd, sendmail, snmp, canna.

S15httpd
S40smartd
S80sendmail
S50snmpd
S50snmpdtrapd
S90canna

© SANS Institute 2005, Author retains full rights.

Control Access

There are a variety of ways to authenticate and restrict access to accounts and resources on the individual computer. There is password control, file control and controlling network access.

Password Control

One of the first things that can be done to secure a computer is to make sure password control is in place. Depending on the operating system, there are up to three different ways to password protect a computer.

With some operating systems, one or more of these methods of authentication can be bypassed completely. Most home computers are set up this way and no password is necessary to operate the computer. In the situation where no authentication is necessary, the whole system may be viewed and controlled with physical access, or even accessed from the network. If the computer is a laptop, it is even more susceptible to being stolen, and if that happens, anything that has been done or recorded on the computer is vulnerable to the thief.

Axiom of Least

When creating administrative passwords, such as the BIOS or bootloader password, give it only to the bare minimum number of people. Don't write the password down, and if it is truly necessary to do so, keep it locked up. Reducing the password to exposure will reduce the possibility of compromise.

When creating user accounts, make sure to give each one the least amount of privileges it needs to accomplish its tasks. In the case of a home user, not all family members need to have administrator or root privileges to be able to install software. This could prevent someone from installing software incorrectly, or prevent inadvertent corruption of the system by malware sent from an already infected computer. Reducing the amount of authority given on the computer reduces the risk of problems occurring.

A general rule is to give the least amount of privilege required to get the job done. There is no point to building a fence around your house if all the windows are left open and doors unlocked.

BIOS Password

The first way to password protect your computer is to go into the system BIOS and set the boot password. When the power is turned on to the computer, the BIOS chip initializes the rest of the computer system to an initial working state. If the BIOS password is set, the initialization process will not complete until the correct password is given.

Boot loader Password

The second way to password protect a computer is limited to the type of boot loader you are

using. The Linux boot loaders, lilo and grub can use a "boot loader" password.

When the BIOS finishes initializing the computer system, it calls a boot loader to load the operating system into memory. When using a boot loader password, these programs will halt the booting process until the correct password has been given.

The method the lilo boot loader uses to configure the password is not very secure. To set the password, just add the following line to `/etc/lilo.conf`:

```
password=<password>
```

The password above is in clear text for all to see. Anyone who can log into the computer can read the lilo configuration file and reboot the machine and enter the boot loader password.

The grub package, depending on the version, comes with a utility to create an encrypted password. The utility prints the encrypted password to the screen. It can then be copied and pasted into the `/boot/grub/grub.conf` file:

```
password -md5 <encrypted-password>
```

There are examples on the net that talk about using lilo and grub bootloader passwords such as:

www.redhat.com/docs/manuals/linux/RHL-8.0-Manual/security-guide/s1-wstation-boot-sec.html

www.linuxforum.com/redhat-security/s1-wstation-boot-sec.html

www.redhat.com/docs/manuals/enterprise/RHEL-3-Manual/security-guide/s1-wstation-boot-sec.html

For RedHat, and the following for Debian:

www.linuxsecurity.com/resource_files/host_security/securing-debian-howto/ch4.en.html

These web sites give instructions on lilo and grub and all provide almost exactly the same instruction for setting a boot loader password.

OS username and password

This password mechanism is the one most people are familiar with. After the operating system has been loaded, it asks for a username and password before the computer can be used.

One best practice is to create accounts for each person who will be using the computer. Allow each one to choose a password. Encourage the use of "strong passwords", which should contain at least 3 of the following - small letters, capital letters, numbers and special characters. Some standard limitations can be made to the way passwords are managed. Typically, regular password expiration, minimum length, account lockout policy and complexity can be managed. Also, some operating systems can keep track of password history - and for the length of the history not allow reuse of old passwords.

There are reasons for each of the above rules to be in place. Regular password expiration can

prevent someone who has harvested passwords on a system from accessing data once the password is changed. A minimum length rule for passwords can increase the difficulty of possible password cracking.

An account lockout policy can prevent or seriously hinder a method such as a “dictionary attack” to attempt to access an account on the system. A dictionary attack is a method an attacker will use which will scroll through each entry in an online text dictionary as a password for a specific user, until it gains access to the computer. With this type of attack, an attacker may attempt to do a dictionary attack or more complex password guessing scheme, to guess the password of an account hundreds of times a minute. A system that has a account lockout policy will lock a user from attempting to log in for a set period of time after a number of bad login attempts. If a lockout policy locks out a user out of the system for fifteen minutes after five bad login attempts, this will frustrate the attempts of a high speed attacker. If a successful high speed attack takes a few days with a few hundred attempts or guesses at a password per minute, limiting the guessing to 5 guesses every 15 minutes would make this method so inefficient as to not be a viable attack method.

Password complexity can hinder someone from using a high speed attack. Password complexity can increase the difficulty of this method of attack, to the point of not being feasible. If each position in a password only had 4 possibilities, such as a, b, c and d, the password would be easily guessable. Increasing each position to 26 possible characters, or the alphabet, increases the complexity. If we increase the possibilities to add numbers, special characters and both upper and lower case numbers, the number becomes much higher and guessing such a password much more complex requiring a much longer time to guess. A good password complexity could require that a password contain three of the following 4 items.

Upper case
Lower case
Special character
Number

Some good complex passwords could include:

Sx1\$bin%
Qer^xana
!lbenfin
4pUblaF8

These are good examples, but since they are now published, should be used as guidelines, not as actual passwords on the system.

Keeping a password history can prevent a user from using the same one over and over each time it expires. It can keep a user from using up to the last X number of passwords. This can prevent an attacker that may have harvested passwords some time in the recent past from using those old passwords to access your system.

Windows XP Pro policy settings

The above OS password rules can be accomplished as follows:

Open up the Control Panel -> Admin Tools -> Local Security Policy -> Account Policies -> Password Policy. Figure 1 Password Policy

Set Enforce Password history for 5

Set Max Age to 180 days (some companies have a rule of 90 days).

Set Password Length to 8 characters.

Enable "Must meet complexity requirements".




Figure 1 Password Policy

Go to Account Policies -> Lockout Policies. Figure 2 Account Lockout Policies

Set to 5 bad attempts.

Set Account Lockout duration to 15 minute lockout for each.




Figure 2 Account Lockout Policy

Unix password policy settings

Solaris, Linux and Irix use System V style password management.⁴ Two files used to control password restrictions are the /etc/shadow file and either /etc/default/login if using Solaris, or /etc/default.defs if using Linux.

Each line in the shadow file follows the following configuration:

```
<username>:<encryptedpasswd>:<lastchngd>:<min>:<max>:<warn>:<inactive>:<expire>:
```

The first two fields contain the user's login name and the user's encrypted password. The third field indicates the number of days since January 1, 1970 that the password was changed. The fourth field indicates the minimum number of days before the password can be changed. The fifth field indicates the number of days after which the password must be changed. The sixth field indicates the number of days before the system will start warning the user to change the password. The seventh field indicates the number of days after the password has expired that the account will become inactive/disabled. The eighth field indicates the actual date the account will expire.

For more information on the password file or the shadow file, type "man passwd" or "man shadow" at the *nix command prompt.

Solaris specific password policy settings

In Solaris, the following command will give the password attributes of all users on the system, including set restrictions.

```
# passwd -sa
```

The /etc/default/login file contains information the system uses to control the login process. In Solaris 9, variables that will control the account lockout policy in the /etc/default/login file include, SLEEPTIME, RETRIES, DISABLETIME and SYSLOG_FAILED_LOGINS. The SLEEPTIME variable controls how many seconds until "login incorrect" message appears after a failed login. The RETRIES variable indicates the number of failed logins that will be allowed before login exits. The DISABLETIME⁵ variable, if greater than zero, controls the number of seconds before PAM exits and login prompt is restored. The SYSLOG_FAILED_LOGINS variable indicates the number of unsuccessful login attempts will take place before it is logged in the system. If this variable is set to zero, all failed login attempts will be logged. Examples of how these variables could be set are:

```
# Sleptime - in seconds
SLEEPTIME=4
# Number of retries allowed - before login exits
RETRIES=5
```

⁴ Frisch, Aileen, Essential System Administration, Second Edition. O'Reilly, 1995: 156

⁵ Man page for login(1) on the Sun Product Documentation website: docs.sun.com/db/doc/816-0210/6m6nb7mdo?a=view

```
# Number of seconds after login exits
DISABLETIME=300
# Number of login failed attempts before the system will start logging the bad attempt
SYSLOG_FAILED_LOGINS=5
```

Linux specific password policy settings

Although it is possible to use the “passwd -S <username>” command in linux to view password attributes of users on the system, it does not give useful information.

Like Solaris, Linux uses a configuration file to control login attributes of its users. This file is /etc/login.defs⁶. Variables that control how the login process works include FAIL_DELAY, LOGIN_RETRIES, LOGIN_TIMEOUT and FAILLOG_ENAB. The FAIL_DELAY variable acts the same as the SLEEPTIME above. The LOGIN_RETRIES works the same was as the RETRIES variable above. The LOGIN_TIMEOUT works the same way as the DISABLETIME variable above. The FAIL_DELAY variable instructs the syslog daemon to log failure of attempts to log in.

```
# in seconds
FAIL_DELAY 4
# Number of retries allowed - before login exits
LOGIN_RETRIES 5
# Number of seconds after login exits
LOGIN_TIMEOUT 300
# Enable logging of failed attempts to log in?
FAILLOG_ENAB YES
```

Other *nix password issues

Although most *nix do not offer a facility for keeping password history, AIX and Digital Unix offer a way to do so⁷. For more information on how to manage password restrictions in *nix, please read page 155 – 163 in Eelen Frisch’s Essential System Administration by O’Reilly and Associates.

File Access

It is possible to limit access to files on the computer. File systems give the ability to control who has what kind of access. This can protect system files, configuration files, sensitive data, and provide a way for users to have private data.

Windows file permissions

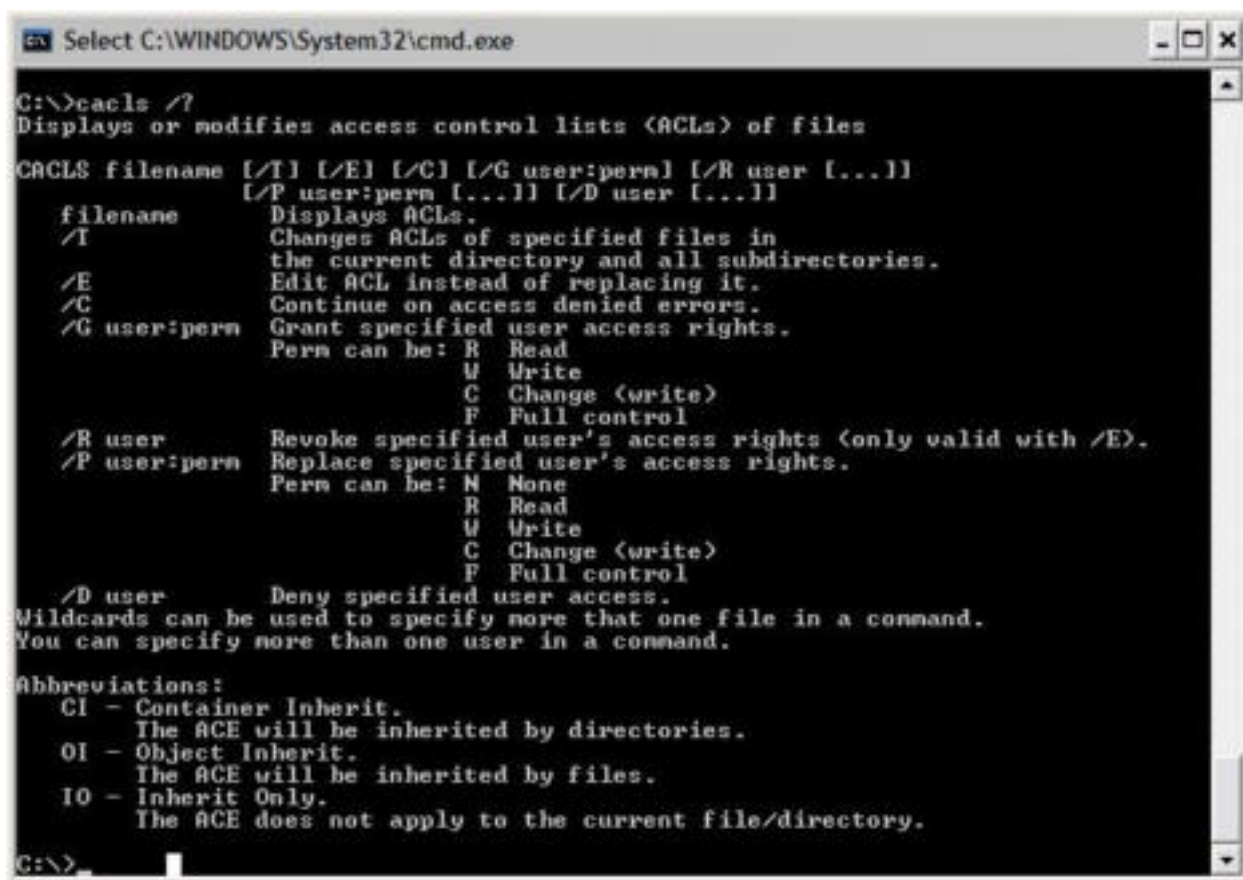
The FAT based file systems, used by Windows 95 and 98 do not have security in mind. Any user can access any file. The NTFS file system, available in Windows NT, 2000, XP and 2003, was designed with security in mind.

⁶ Man page for login(5) on the Linux Valley mirror of The Linux Documentation Project website: www.linuxvalley.it/encyclopedia/ldp/manpage/man5/login.defs.5.php

⁷ Frisch, Eelen, Essential System Administration, Second Edition. O’Reilly, 1995: 217

With an OS that uses the NTFS based file system, a security token is handed to the user upon login. Windows uses this security token to determine access rights when an ACL is in place⁸.

In Windows XP Home, the only method to view access control information is via the `cacls.exe` command. The only information in Windows XP Home for this command is to type “`cacls /?`”, see Figure 3 `Cacls.exe`.



```

Select C:\WINDOWS\System32\cmd.exe

C:\>cacls /?
Displays or modifies access control lists (ACLs) of files

CACLS filename [/I] [/E] [/C] [/G user:perm] [/R user [...]]
               [/P user:perm [...]] [/D user [...]]
  filename      Displays ACLs.
  /I            Changes ACLs of specified files in
               the current directory and all subdirectories.
  /E            Edit ACL instead of replacing it.
  /C            Continue on access denied errors.
  /G user:perm  Grant specified user access rights.
               Perm can be: R Read
                           W Write
                           C Change (write)
                           F Full control
  /R user       Revoke specified user's access rights (only valid with /E).
  /P user:perm  Replace specified user's access rights.
               Perm can be: N None
                           R Read
                           W Write
                           C Change (write)
                           F Full control
  /D user       Deny specified user access.
Wildcards can be used to specify more than one file in a command.
You can specify more than one user in a command.

Abbreviations:
  CI - Container Inherit.
      The ACE will be inherited by directories.
  OI - Object Inherit.
      The ACE will be inherited by files.
  IO - Inherit Only.
      The ACE does not apply to the current file/directory.

C:\>
```

Figure 3 `Cacls.exe`

With Windows XP Pro, the files can be controlled with DACLs or Discretionary Access Control Lists. To make sure this option is available, in a file explorer window go to the tools menu, click on “Folder Options”, click on the “View” tab, and deselect the “use simple file sharing” box, then hit the “Apply to All Folders” button. Now it is possible to right click on the folder or file you want to control access to, go down to Properties and click on it, then click on the Security tab. See Figure 4 DACLS below.

⁸ Bradford, Ed and Mauget, Lou, Linux and Windows Interoperability Guide. Prentice Hall PTR, 2002: 365 - 369

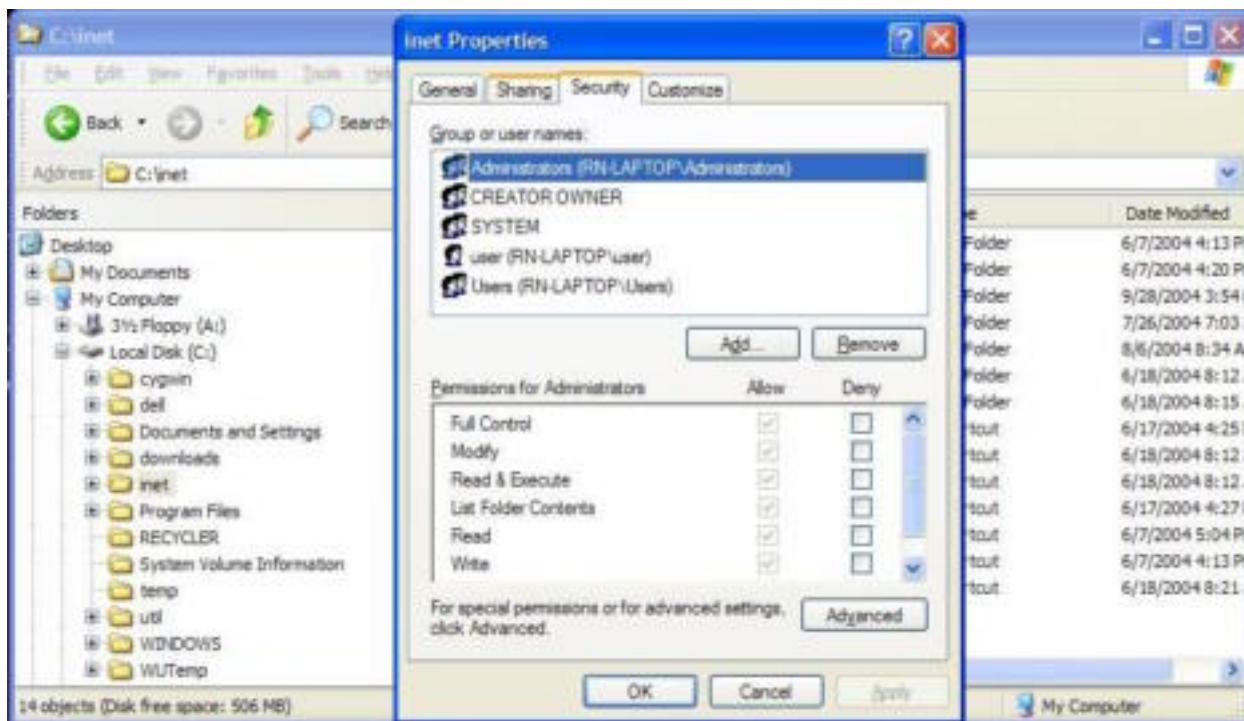


Figure 4 DACLS

More information on controlling access in Windows to files can be found at the following web address. www.microsoft.com/windowsxp/using/security/learnmore/accesscontrol.msp

Unix file permissions

There are three categories of file permissions on *nix systems. They are user, group and other. Within those categories, there are three types of permissions a file may have, which are read (r), write(w) and execute(x). For instance, the user category indicates that only the user, or root may perform the task of read, write and/or execute.

When a file listing is done, the permissions of the file are given in the following order:

`-rwxrwxrwx`

The `rwx` category on the right indicates that the file is readable, writeable and executable by all other users on the system. The middle `rwx` indicates that the file is readable, writeable and executable only by the group that owns it. The `rwx` on the left indicates the permissions for the owner of the file - as above - read, write and execute.

Executing the “`ls -l`” command on a file will list the file and its properties. An example from Kernighan and Pike’s The Unix Programming Environment⁹:

⁹ Kernighan, Brian and Pike, Rob, The UNIX Programming Environment, Prentice-Hall, 1984: 53

```
$ ls -l /etc/passwd
-rw-r--r-- 1 root 5115 Aug 30 10:40 /etc/passwd
$ ls -lg /etc/passwd
-rw-r--r-- 1 adm 5115 Aug 30 10:40 /etc/passwd
```

The listing above indicates that the password file is owned by user “root” and by group “adm”, is 5115 bytes and was last modified on August 30th at 10:40am. It also indicates that group “adm” can only read the file, and the “other” field indicates that all others, on the system can only read the file as well. Another example from Kernighan and Pike¹⁰:

```
$ ls -l /bin/who
-rwxrwxr-x 1 root 6348 Mar 29 1983 /bin/who
```

This set of permissions indicates that the user, root, has read-write-execute permissions, and so does the owning group. The permissions for all others are read and execute only.

To change permissions on a file, use the `chmod` command. This command will change the mode of the file. The format of this command is:

```
$ chmod options permissions filename
```

The permissions can be expressed in a numerical fashion, in other words, 4 for read, 2 for write and 1 for execute. Add the numbers together that are needed for the category and that is the number to use.

An example of the numbers that can be used is:

```
7=rwx
6=rw-
5=r-x
4=r--
3=-wx
2=-w-
1=-x-
0=---
```

In other words, to give the permissions above for the `who` command, one would execute the following command.

```
$ chmod 775 /bin/who
```

If I wanted to create a file `/home/roy/newFile` and give it permissions for `-rwxr-x---`, or read/write/execute for me, read/execute for the group and no permissions for others, I would execute the following command:

```
$ chmod 750 /home/roy/newfile
```

Also in *nix there is a command called the `umask`. This is a command that will set the default

¹⁰ Kernighan, Brian and Pike, Rob, The UNIX Programming Environment. Prentice-Hall, 1984: 55

permissions for newly created files. This is usually done in login scripts so the files created after a user logs in have a desired level of security. The `umask` command uses numerical values similar to the ones used with the `chmod` command. To get the appropriate value, figure out the permissions that are needed, then subtract from 777. To get the above permissions by default for the `newfile` and any other new files created, set the `umask` to 027 by putting the following line in your `.login` file.

```
umask 027
```

More information on the `umask` command can be found by typing `man umask` at the command prompt.

Controlling network access

Different operating systems use different, although similar, methods to control access to the computer from the network. The main way to do this is to use a firewall. Unix has a couple more concerns, first with `tcpwrappers`, second with the “R” services.

Firewall

Think of a computer connection to the network as a sieve, each hole open to network traffic of a specific type. Web traffic goes through one, FTP traffic through another, Telnet through another, ssh through another, and so on, for more than sixty five thousand holes or ports. A firewall will block all ports except those that are specified in a “firewall rule set”.

An advantage of a firewall is to block unwanted random scanning of ports on the computer by infected hosts on the network that are attempting to exploit vulnerabilities and/or spread infection. A firewall could also block attacks that target your specific computer. There is a drawback to firewalls however. In any server-client relationship such as the world wide web, citrix, ftp, ssh, and other services, both client and server need to open holes to be able to communicate. If one or the other does not have the appropriate port or hole open, then there can be no communication.

Windows Firewalls

There are several sites on the internet that review Windows firewalls. One such site is:

www.computerproblems.org/forums/showthread.php?t=333


Two of the popular Windows firewalls are ZoneAlarm and BlackIce. BlackIce is part IDS, or Intrusion Detection System. It keeps track of packets that hit the firewall and by default, alerts the user to possible intrusions, as well as displaying attempts to connect to the computer. Zone alarm on the other hand does not function as an IDS. It can, however, not only block traffic coming into the computer, but also has the facility to control outgoing packets as well.

One main feature of BlackIce is that it keeps track of and displays attempts to connect to the

firewall. ZoneAlarm does not do this. A main feature of ZoneAlarm is the ability to block and manage packets leaving the computer as well as incoming packets, which BlackIce does not possess.

I have chosen to use the BlackIce firewall as an example Windows firewall. Default installed settings will work well for most sites.

To view BlackIce settings, double click on the BlackIce icon in the Windows toolbar.

 The icon is indicated by the arrow in the above picture and the following window should appear - Figure 5 BlackIce.




Figure 5 BlackIce

To add a firewall rule go to the Tools menu, Figure 6 BlackIce Tool Bar, go down to “Advanced Firewall Settings”, and the following window, Figure 7, will pop up with the default firewall rule set.

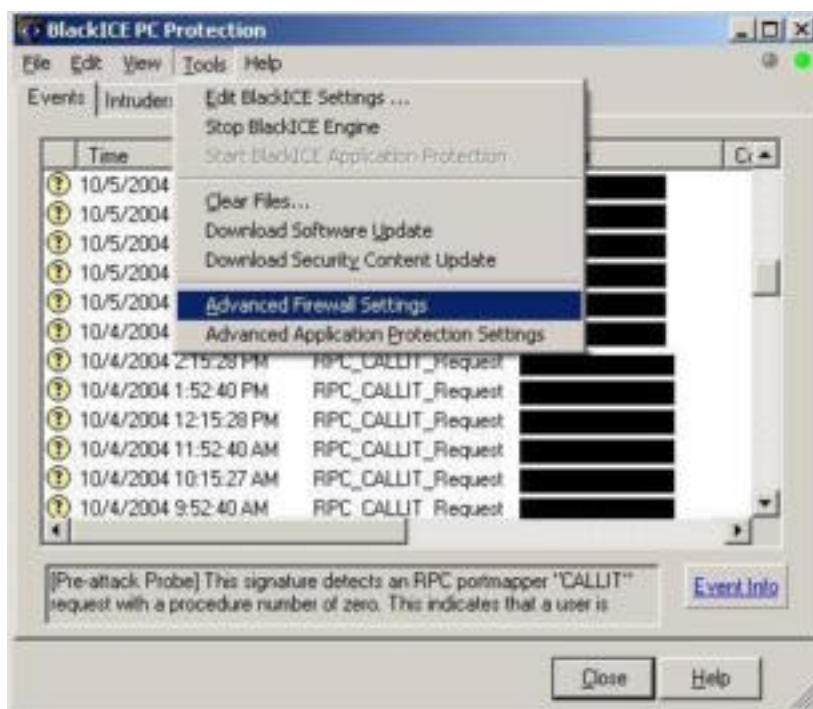


Figure 6 BlackIce Tool Bar

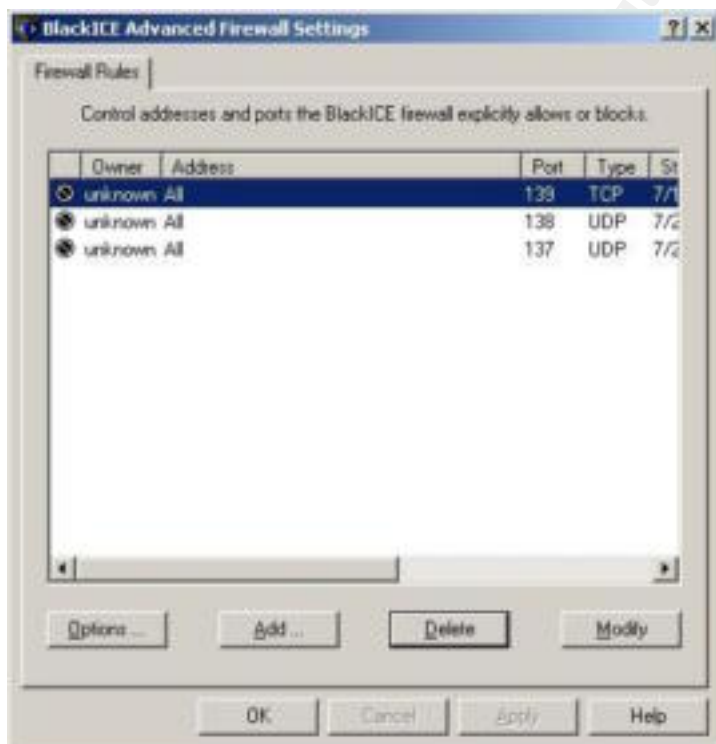


Figure 7 BlackIce Firewall Rules

To add a rule, hit the “Add” button, and the following window will pop up, Figure 8 BlackIce Add Rule.

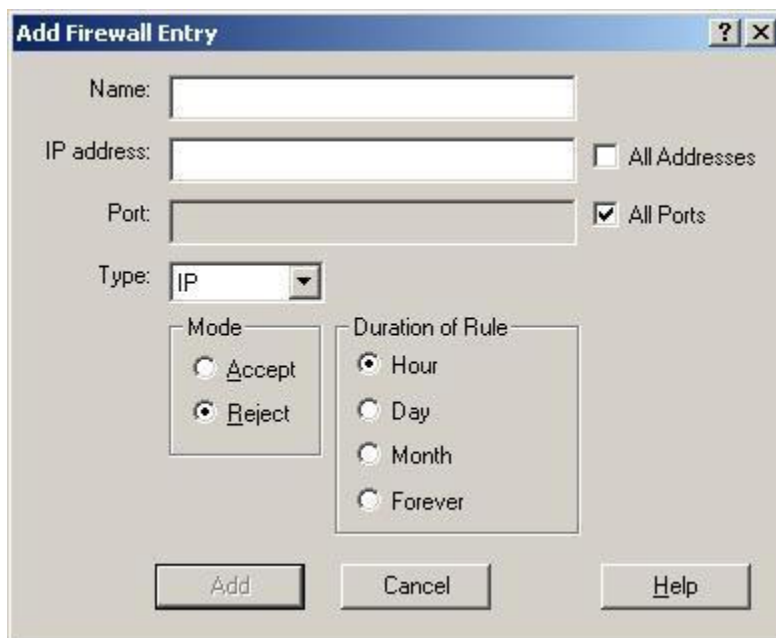


Figure 8 BlackIce Add Rule

An example of a rule that rejects tcp snmp traffic would be Figure 9 Snmp Block:

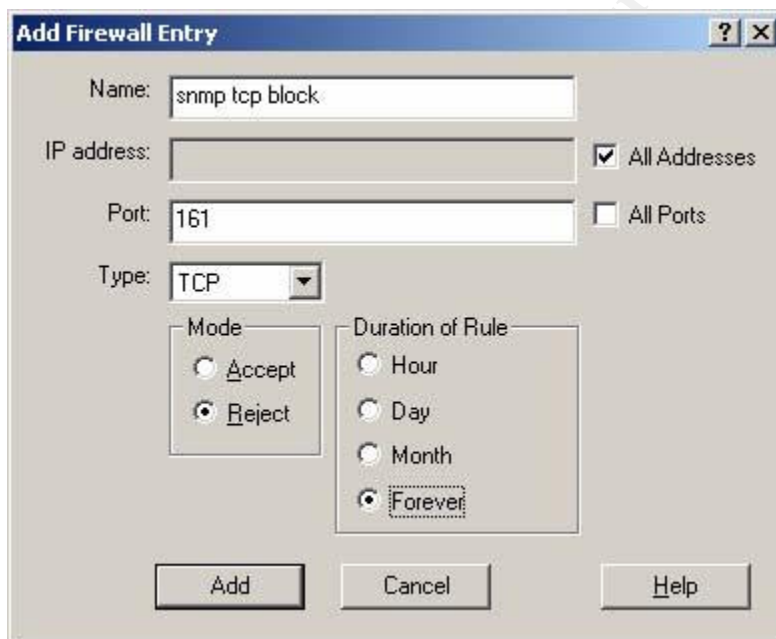


Figure 9 Snmp Block

Make sure to hit the “Add” button to apply the firewall rule.

An example of a firewall rule to allow all traffic from a specific IP address for a limited time –
Figure 10 Temporary Host Connection:

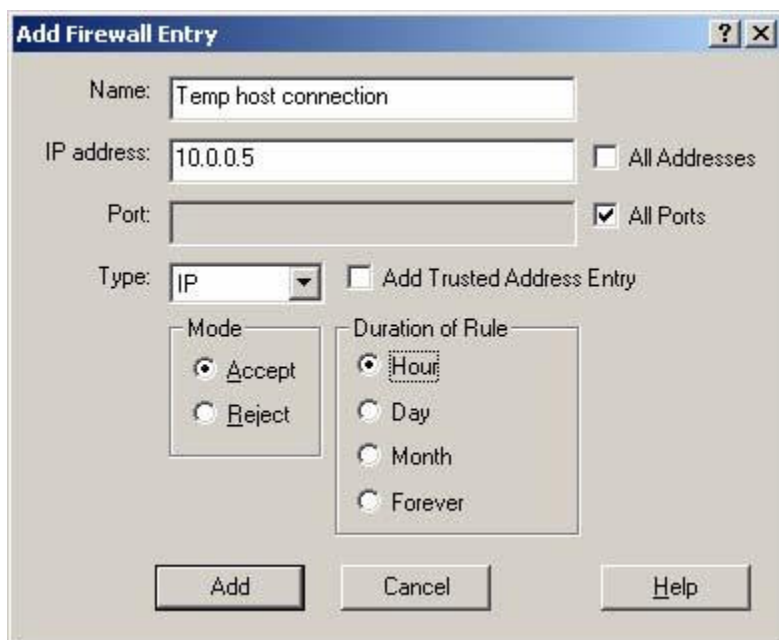


Figure 10 Temporary Host Connection

It is appropriate to occasionally trust a single machine for all ports and protocols. By selecting the “Add Trusted Address Entry”, this can be accomplished. Any traffic from this machine will be trusted, no matter what it is, and no alerts will be given or logged in this case. See Figure 11 Trusted Gateway for an example.

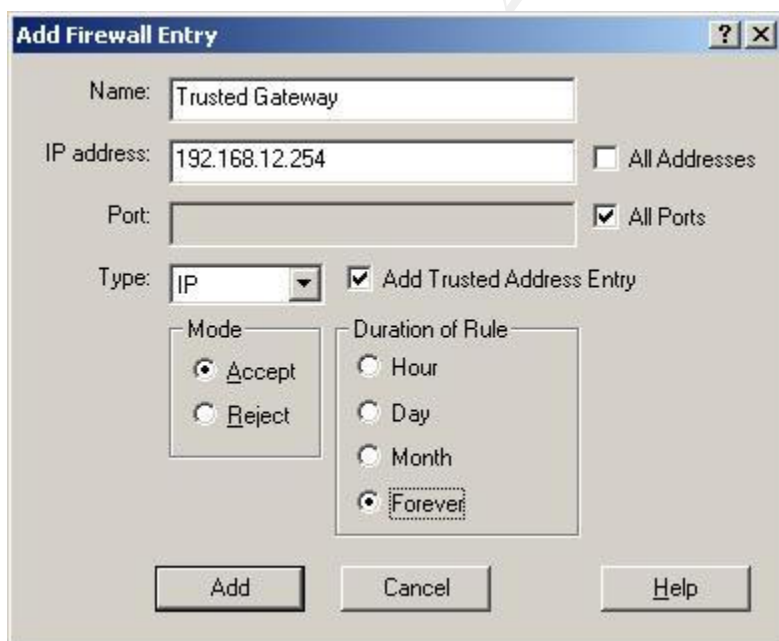


Figure 11 Trusted Gateway

Again, when finished filling out the options for the rules that are needed, make sure to hit the “Add” button to add the rule to the firewall rule set.

Solaris firewall

Solaris currently doesn’t come with a firewall. Solaris 10 will be released with the *nix firewall ipfilter. This firewall is available for different versions of *nix and can be found at coombs.anu.edu.au/~avalon/.

The firewall can either be built from source, or the web page above has a link to pre-built packages that can be installed with “`pkgadd`” command. For more information on how to use `pkgadd`, type “`man pkgadd`” at the command line.

The ipfilter firewall rules can get very complicated, to the point of allowing only a single port on a particular interface from a specific machine. The best place to find documentation on ipfilters is:

www.obfuscation.org/ipf/

Some common rules that are good for any ipfilters firewall are:

```
### basic /etc/opt/ipf/ipf.conf firewall ruleset
# first thing is to write the rules for “allowing” all traffic
# from the local subnet
pass in quick from <local-subnet> to any keep state

# allow tcp http traffic in, on interface hme0 from machine 10.0.0.5,
# allowing returning/answering traffic to be allowed in
pass in quick on hme0 proto http from 10.0.0.1 to any keep state
# or
pass in quick on hme0 from 10.0.0.1 to any port = 80 keep state

# traffic that should not be allowed on a routeable network
block in quick from 192.168.0.0/16 to any
block in quick from 172.16.0.0/16 to any
block in quick from 10.0.0.0/8 to any

# possible spoofed or garbage traffic
block in quick from 0.0.0.0/8 to any
block in quick from 127.0.0.0/16 to any

# Allow pings to local corporate network
pass out quick proto icmp from any to any keep state
pass in quick proto icmp from <IP-range> to any icmp-type echo
pass in quick proto icmp from <IP-range> to any icmp-type echorep

# allow all outbound traffic, and returning/answering traffic to be allowed in
pass out quick proto tcp from any to any keep state
pass out quick proto udp from any to any keep state

# Block all inbound traffic that hasn’t been specifically allowed
## return tcp reset for tcp traffic
block return-rst in quick proto tcp all
## for udp and icmp traffic, block without responding
```

block in all

Linux firewalls

Three available firewalls for Linux are ipfilters, ipchains and iptable. The ipfilters firewall rules work in Linux just as they do for Solaris. Since the ipfilters program is portable across several *nix platforms, the syntax for the firewall rule set is the same as with Solaris.

The ipfilters web page does not list Linux as a supported platform. However, in the HISTORY file that is part of the Ipfilters source, it states that ipfilters was ported to Linux in 1997, and has had updates to Linux specific code as recent as May 2004. A list of mirrors where ipfilters source code is available at: coombs.anu.edu.au/~avalon/#Mirrors.

Ipchains and iptables are native to Linux. Ipchains is a predecessor to iptables, and it does not do stateful packet filtering. Stateful packet filtering is basically keeping track of connections, and allowing those that might change ports to continue without being blocked. One example of a protocol that doesn't stay with a single port is ftp.

Chapter two in the Linux Security Cookbook¹¹ is a valuable introduction and reference for understanding and writing firewall rules for both ipchains and iptables.

I prefer iptables, because of the stateful nature of the firewall. Some example iptables firewall rules can be found below.

```
#####
## basic /etc/sysconfig/iptables rule set
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
:ETH0-IN - [0:0]

#####
#### INPUT Rules ####
##### ALLOW connections that are already established or related... #####
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT

#####
# allow all traffic from the local subnet
-A INPUT -s <local-subnet> -j ACCEPT

#####
# allow tcp http traffic in, on interface eth0 from machine 10.0.0.5,
# allowing returning/answering traffic to be allowed in
-A INPUT -I eth0 -m tcp -p tcp -s 10.0.0.5 --dport 80 --syn -j ACCEPT

#####
#### ALLOW SSH ####
-A INPUT -p tcp -m tcp --dport 22 --syn -j ACCEPT
-A INPUT -p udp -m udp --dport 22 -j ACCEPT
```

¹¹ Barrett, Daniel J., Silverman, Richard E. and Byrnes, Robert G., Linux Security Cookbook. O'Reilly, 2003: 23 - 48

```
#####
#### ALLOW PINGS ####
-A INPUT -p icmp -m icmp --icmp-type echo-reply -j ACCEPT
-A INPUT -p icmp -m icmp --icmp-type echo-request -j ACCEPT
#### END INPUT RULES ####

#####
#### REJECT EVERYTHING ELSE ####
-A INPUT -p tcp -m tcp --syn -j REJECT
-A INPUT -p udp -m udp -j REJECT

COMMIT
```

Unix Tcprappers

This is done by tcpd, a daemon that “wraps” services giving access only to specific machines or networks. It does this for services that are run by the meta-server above. Some services that can be wrapped are ssh, ftp and lpd. The rules that determine which machines or networks have access to these services can be found in /etc/hosts.allow and /etc/hosts.deny.

The hosts.deny file should have the only the following line, uncommented:

```
ALL:ALL
```

This will block everything that is not explicitly allowed in the hosts.allow file. The purpose of this file is to contain rules to allow traffic through and are much simpler than *nix firewall rules. The first column should be the name of the service or daemon being wrapped, followed by a colon, then the computer or network to be allowed through. The following rule is an example of the wrapper allowing ssh traffic in from the internal network address 10.0.0.1 – 10.0.0.255.

```
sshd: 10.0.0.
```

For more information on tcpd, hosts.allow and hosts.deny use the man page, like “man tcpd”, “man hosts.allow” or “man hosts.deny”. Instruction provided by these man pages can guide one through configuring the tcprappers services.

Unix “R” services

Another important unix computer vulnerability can be found with the remote or “R” services such as rsh, rexec and others. Although these services can be wrapped with tcpd, they are still dangerous, because they can be configured to bypass authentication altogether. This could give someone with malicious intent easy access to the computer system. It is important to disable these in inetd and xinetd configuration files.

The only time using the “R” services might be ok, is if they work in an un-routable network, with no ability to allow outgoing or incoming traffic to this network. “R” services are dangerous, and anyone using them in an uncontrolled environment is leaving their doors unlocked, and windows open.

Logging

When a program fails, how do you know what went wrong? Can the operating system keep track of what is going on and log errors and anomalies? The web server is acting strange, is there a way to find out what is going on? All operating systems have logging facilities and many programs, servers and services also do as well.

*nix Logging

BSD style logging occurs with the “syslogd” daemon. The file used for configuring this daemon is `/etc/syslog.conf`. The messages are sent by the syslog daemon to a variety of locations. The first column in the `syslog.conf` file indicates the type of message to send, and the second column is the target for messages to go. The default settings are useful for most computers. If you want to change the default settings, check the man page for `syslog.conf` for directions for configuration of the `syslog.conf` file.

Most useful messages are sent to the `/var/adm/messages` on Solaris, and `/var/log/messages` in Linux. Look the `/var/adm` or `/var/log` directory for other useful logs from which to glean helpful information.

Network printers and other network devices often come with syslog capabilities, and can send their logs over the network to a syslog server. Again, check the man page on `syslogd` to configure the machine to catch and log these networked messages. The `syslog` startup script may need to be modified to enable this.

Windows logging

Windows XP Home logging can be set up from the “Event Viewer” in the “Administrative Tools”, found in the Control Panel. There are three categories of logging which are application, security and system logging, as seen in figure 12. Right clicking on the “Application” icon, go down to “Properties”, and a window will show the properties of the log file, as in Figure 13. Default settings can be found in Figures 13 Application Properties & Figure 14 Application Filter.

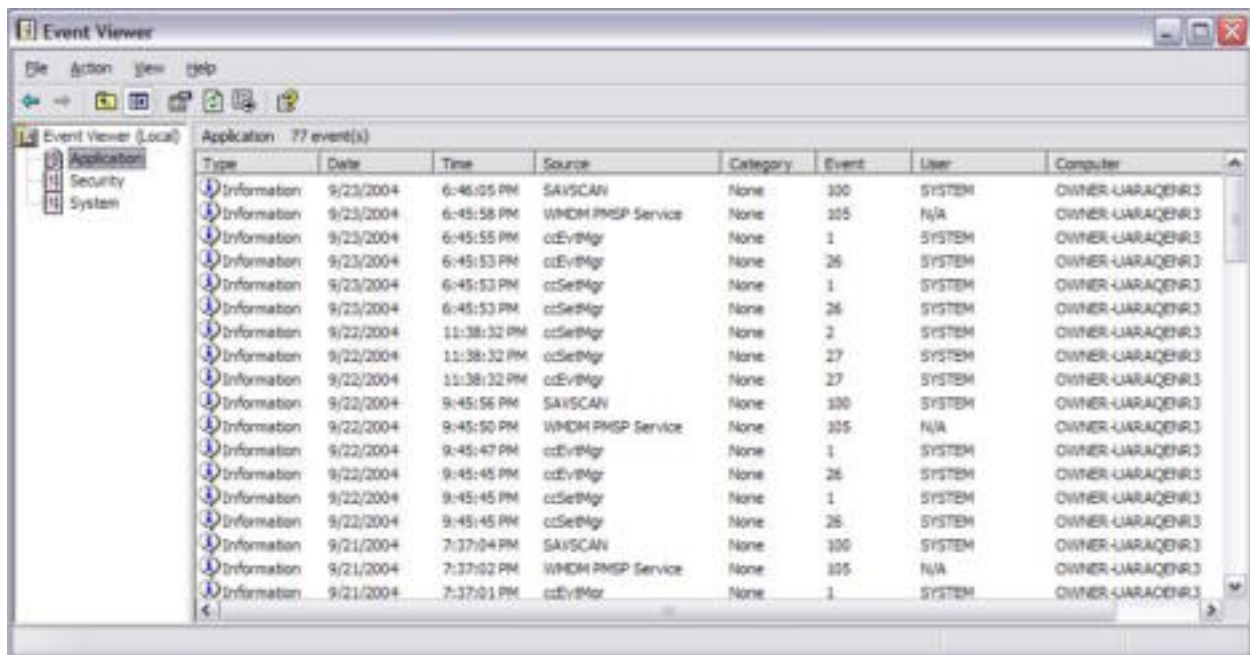


Figure 12 Application Log

Double click on an event in the right hand pane, as in Figure 12 to see the properties of the event in question. For more information on how to manage the logs, and even to use them to debug system problems, right click on the right hand pane in Figure 12 and go to help.

© SANS Institute 2005, All rights reserved.




Figure 13 Application Properties

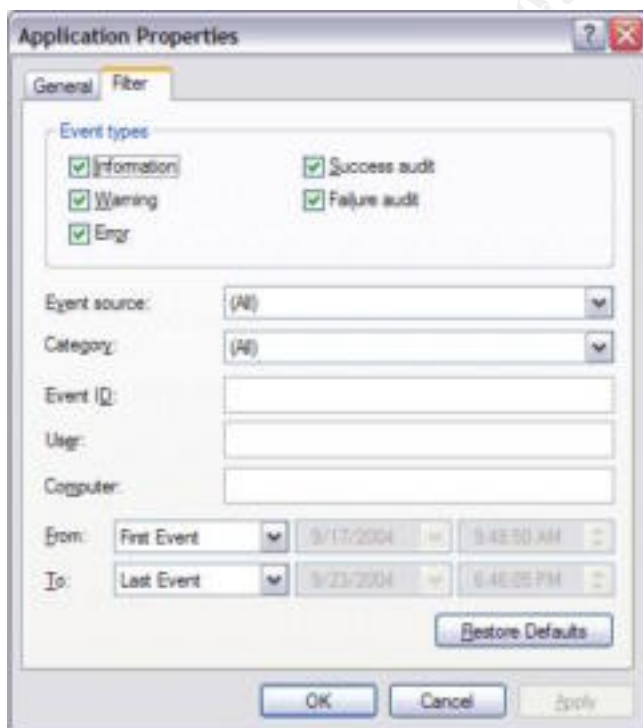


Figure 14 Application Filter

Windows XP Pro has more facilities for logging. In the “Local Security Policies” in the “Administrative Tools” found in the “Control Panel”, under the Local Policies -> Audit policy, several things may be tracked. Some good things to track for both Success and Failure are as shown in Figure 15 Audit Policy below:



Figure 15 Audit Policy

The setting for “directory service access” is not enabled to Success and Failure because it is not connected to an active directory domain.

Firewall logging

Firewalls have facilities for logging activity that passes through them, or is blocked by it. Often, even though they have the ability to log, firewalls are not configured to do so.

Be cautious with the logging configuration. If limits are not set on the size of the logs, an attacker can fill up your file system by just attempting to communicate with your computer millions of times in a short period of time. This could cause the OS to crash.

Windows BlackIce logging

Keep in mind that attempting to read the BlackIce logs can be difficult and require a trace file decoder to make sense of them. Many can be found on the net. One such program is Ethereal, which is available at www.ethereal.com, which is available for Windows, Solaris and Linux.

BlackIce has a GUI interface for selecting logging options. As above, click on the BlackIce icon in the menu bar to start the GUI interface. Go to the “Tools” menu, and click on “Edit BlackICE Settings ...” as seen in Figure 16 Edit BlackICE Settings.

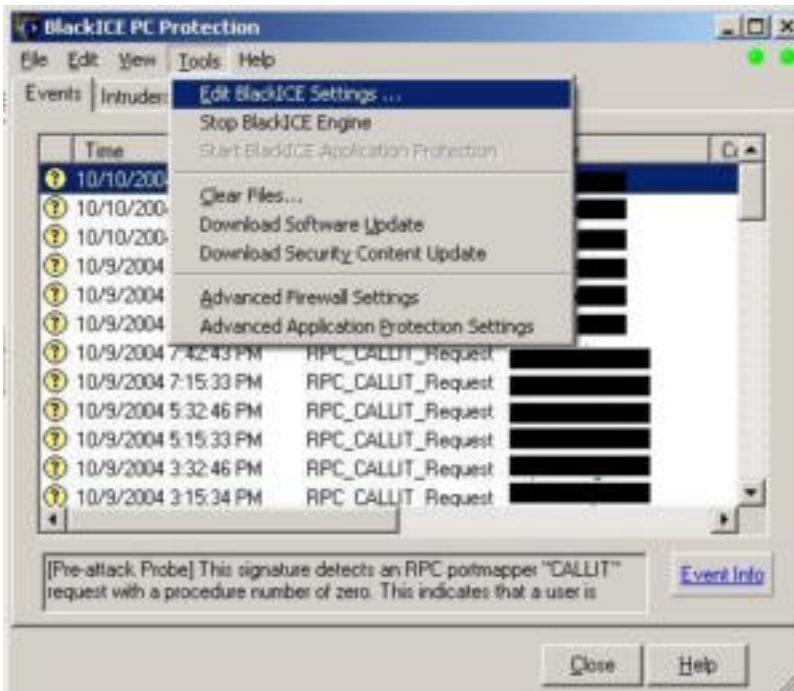


Figure 16 Edit BlackIce Settings

Click on the “Packet Log” tab, which will show Figure 17 Packet Log.



Figure 17 Packet Log

If the Help button is clicked on the above window, the following text is available - plus further instructions on the options on the tab above:

“You can record all the network traffic that passes through your computer. BlackICE generates files, called packet logs that contain detailed information about an intruder's activity.”

Figure 18 shows the settings for Evidence Logging.



Figure 18 Evidence Log

The following text is available if the “Help” button is clicked on the above window:

“You can capture the specific packet that set off a protection response and save that packet in a special file. If you need to, you can use the evidence in this file for an investigation.”

The log files can be found in the root installation directory, which by default is: `c:\Program Files\ISS\BlackICE`. The files should be named `evdXXX.enc` and `logXXX.enc`, where the XXX is the number assigned by BlackIce. Once the current log fills up, one is added to and a new file written. When the number gets to the limit set in the configuration below, the numbering scheme starts back at zero and as the log fills up, increments again.

Linux iptables logging

In Linux, the iptables firewall passes information to the syslog daemon, which does the actual logging. No configuration of the syslog daemon needs to take place in order to enable the syslog daemon to log iptables messages. The output of the iptables logging is sent by the syslog daemon directly to `/var/log/messages` by default. With iptables, in order to log a specific behavior, like logging activity on port 80 if you have a web server, two rules need to be written. One rule needs to be made to log the activity, and one to accept or reject the activity. An example:

```
-A ETH0-IN -i eth0 -p tcp -m tcp --dport 80 --syn -j LOG
-A ETH0-IN -i eth0 -p tcp -m tcp --dport 80 --syn -j ACCEPT
```

The logging facility and options for logging with iptables can get quite complex. More explanation of how to use the logging facility can be found at:

iptables-tutorial.frozentux.net/iptables-tutorial.html#LOGTARGET

Solaris ipfilter logging

As with Linux, in Solaris, the ipfilter firewall passes the logging information to the syslog daemon, which does the actual logging. However, before logging will start, the syslog daemon needs to be configured to accept the information from the firewall and know where to place the logs. The syslog configuration file, `/etc/syslog.conf` must have the following line, typically at the end of the file:

```
local0.debug                                /var/log/ipflog
```

After this configuration edit, syslogd must be restarted. To do this, kill the syslogd daemon with the HUP signal and it will re-read the `syslog.conf` configuration file. This must be done as root.

```
# ps -ef | grep syslogd
# kill -HUP <pid-of-daemon-found-in-first-column-of-above-output>
```

This will cause the syslog daemon to re-read its configuration file so it will start the logging process. Syslogd will now write messages passed to it by ipfilters to `/var/log/ipflog`.

An example ipfilter firewall rule that includes logging would be:

```
pass in quick log on hme0 proto tcp from any to any port = 80 keep state
```

As mentioned above, after changing ipfilter firewall rules, reload the firewall by doing the following as root:

```
/etc/init.d/ipfboot reload
```

Be cautioned, however, depending on the version of ipfilter that is used, if you do the following on the command line:

```
$ /etc/init.d/ipfboot stop
$ /etc/init.d/ipfboot start
```

instead of using following command,

```
$ /etc/init.d/ipfboot reload,
```

all network connections will be severed.

Server Logging

Some people may choose to run servers on their machines, such as apache, mysql, tomcat, etc. All these servers come with facilities to log access to their services. These servers place a “log” or “logs” directory in the root directory where they are installed. For instance if the apache web server is installed in `/usr/local/apache` on a *nix machine, the log directory can be found in `/usr/local/apache/logs`. There are two standard logs – the `access_log` and the `error_log`, and if configured, apache can log more information based on the optional modules installed with the apache server. The configuration file to configure the way the server works is often found in the “conf” directory in the root of where the server is installed. The configuration file for apache, if it has been installed in `/usr/local/apache` can then be found in `/usr/local/apache/conf`. The level of logging, or what will be logged and if it is turned on is controlled via the conf file.

Tomcat and mysql work the same way. Other servers also have configuration files that not only control different aspects of how the server works, but also how it logs and where the log files are kept.

If the server is observed having problems or throwing errors, the log files are an excellent place to start the debugging process.

Filesystem Integrity Logging

Tripwire

With tripwire, a baseline picture or “snapshot” of the file system can be taken, and when run with the “check” option, it will report if any file has been changed, deleted or added to the file system. Tripwire is available for many operating systems, free for some Linux systems at www.tripwire.org, and a pay-for version is available for Solaris, AIX, HP-UX, FreeBSD, and Windows at: www.tripwire.com/products/servers/platforms.cfm.

All operating systems have a method for scheduling regular tasks. It is good practice that, once a baseline snapshot is created, to run a regular tripwire check, and review the logs to make sure nothing unusual has happened.

If a new baseline isn’t made after installing software, the new software will likely be in the new tripwire report. So each time a new software package is installed, the report will get longer. Making a new baseline tripwire database or snapshot is important to be able to clearly and quickly be able to tell that something is amiss with the computer system.

Logs can get long and tedious if the tripwire hasn’t been properly configured. It is possible to configure it to report at different levels of warning, or selecting directories or files to exclude altogether from the tripwire. One such directory that would be good to exclude is the “cache” kept by the web browser. Another possible directory to exclude is where pictures, papers, music, or other data that might change frequently are kept.

A vulnerability of a tripwire is that an attacker can modify or replace the baseline database or

configuration file. So, known best practice is to put the configuration file and database on a media that is read-only, like a CD-R or USB stick that has a switch than can be made read-only.

To further protect the tripwire configuration and database, the files can be encrypted or hidden using stegonography.

Last, by placing the database and configuration under configuration management, such as RCS, CVS, or a similar change control database to keep track of differences between revisions.

Another product that is similar to the “tripwire” product described above is called Aide – Advanced Intrusion Detection Environment is reputed to have helped the Debian Linux administrators discover the brake-in to their project machines. News of this event can be found at: lists.debian.org/debian-announce/debian-announce-2003/msg00003.html. This product is available for AIX, BSD, Linux, SunOS/Solaris and Posix based operating systems. It can be found at: sourceforge.net/projects/aide.

© SANS Institute 2005, Author retains full rights.

Scanning for Malware

Malware is a term for software that installs itself on a computer without your consent. There are several kinds of malware. Advertisers like to track what users look at on the web. They target advertising to the user specifically - this is called adware. There is a similar kind of software that can find its way onto a computer called spyware. It can track what is done on the computer, and sends the information to a remote server. One typical type of spyware is a “keystroke-logger”, which is a program that records your keystrokes. This software can gather usernames and passwords, credit card numbers, bank account information – anything you type while your computer is on.

Another kind of malware is the well known “computer virus”. Most of the time, all malware is lumped into this category. This software installs itself on the computer, then reproduces itself. It can crash the computer system, and/or attack other computer systems on the network. A trojan is a piece of software that masquerades as a game, e-card, e-mail attachment or other program to try to entice the user to start it up so it can install the embedded malicious code. Due to security holes in some e-mail packages, simply viewing an email can create a new victim.

According to a Reuters article¹², companies lost \$55 billion in 2003 due to viruses. A Yahoo article¹³ states that 20% of Dell’s tech support calls are related to spyware.

What protection is there against malware? Often malware exploits security holes in software products and operating systems. One way is to keep the operating system and software up to date by practicing regular patching, which is covered in the next section.

Besides patching, it is very important to have anti-virus software installed on the computer. Besides scanning e-mail as it comes into the computer for known malware, it can also provide real time protection of the filesystem – keeping the configuration and system files secure. Protection is a very important part of anti-virus software. However, there are many people who don’t buy or install anti-virus software until after they have been infected.

Anti-virus software keeps definitions or signatures of malware. As new malware is released into the wild, the software companies release new definition files that can scan for, clean up and protect from these new viruses. A feature in this software will allow the user to connect to the anti-virus company virus definition server and download new definitions as they are created to deal with new viruses. The update process can be automated or scheduled in most cases.

A site that reviews anti-virus software and technology for Windows, Linux and Netware is:
www.virusbtn.com/vb100/archives/products.xml?table

¹² Reuters, “\$55bn virus damage costs for businesses last year”. www.silicon.com, 2004, www.silicon.com/software/security/0,39024655,39117842,00.htm

¹³ Bridis, Ted “Group, Dell Launch Anti-Spyware Campaign”, Associated Press, story.news.yahoo.com/news?tmpl=story&u=/ap/20041016/ap_on_hi_te/internet_spyware&e=1

Some popular Windows anti-virus software can be found at:

www.kaspersky.com

www.symantec.com

www.mcafee.com

Two popular Windows adware removal software can be found at:

www.lavasoftusa.com

www.spybot.info/en/index.html

Although *nix does not suffer from viruses in the same manner as Windows, it still battles malware. Two free products that help check for problems on *nix based machines are:

www.chkrootkit.org

www.rkhunter.org

© SANS Institute 2005, Author retains full rights.

Patching

As software writers, engineers and companies continue to improve their products, they release patches; in other words, updates and fixes to plug the holes and squash the bugs as well as increasing the feature sets of their products.

One continuing problem for people creating software is the “buffer overflow”. This can be due to the way a person writes code, or how the OS handles memory requests, or how a compiler creates code. When a section of memory is allocated for a variable or set of variables in a program, there is the possibility that data input into these variables can “overflow” the bounds of the allocated memory. When this happens, an attacker can use this to take control of the computer. For instance, he could attack a mail server by sending specially crafted email to the server, which can create an overflow condition on the servers and the attacker can then gain access to the system. An example of this can be found at secunia.com/advisories/8194. The buffer overflow is a common problem, and vendors are continuing to release patches to fix problems like this.

The buffer overflow is just one of many problems or holes which vendors release patches for. So, it is important to follow product development, not only for updated, looked for features, but also for potential security and bug fixes. If the computer is in a production environment, it is important to test the patches before deploying them at the whole site.

Applications

Most often malware will take advantage of bugs or features of programs. For example, with Microsoft Word, there are “Macro viruses”. A macro is a way to record a bunch of commands in Microsoft word to create a shortcut. In 1999, the Melissa Word macro virus hit computers. An article in August of 2000 at www.securityfocus.com/infocus/1278 by Denis Zenkin called “Understanding Macro Viruses” discusses macro viruses, and what can be done to clean up after them, and protect against them¹⁴.

Some applications and servers that have required patches or updates due to security related problems in 2004 are:

wu-ftpd: This is a ftp server, widely used in the Linux world. The advisory: www.ciac.org/ciac/bulletins/o-095.shtml, www.securityfocus.com/bid/9832

mysql: Widely used database, this problem is an authentication bypass issue. The advisory: www.kb.cert.org/vuls/id/184030, www.securityfocus.com/bid/10654

apache: Very popular web server had a vulnerability in an underlying secure protocol in mod_ssl. The advisory: www.kb.cert.org/vuls/id/303448, secunia.com/advisories/12077

¹⁴ Zenkin, Denis “Understanding Macro Viruses”. www.securityfocus.com, 2001, www.securityfocus.com/infocus/1278

Microsoft Word: Popular Microsoft word processing software, possible DOS vulnerability and possible remote system compromise. The advisory: secunia.com/advisories/12758

Information on recent patches, and what they fix can be found in many pages on the net. Some web sites that give information on security holes and appropriate patches can be found at:

www.linuxsecurity.com	www.virusbtn.com/news/latest_news/index.xml
www.securityfocus.com	www.ciac.org/ciac/index.html
www.us-cert.gov	www.cert.org
secunia.com	netsecurity.about.com
www.securityfocus.com	isc.sans.org
www.securiteam.com	www.osvdb.org

Operating Systems

Operating Systems also have problems and provide updates. There are several technologies that provide this service. For Windows there is the Automatic Update service, SMS and SUS. For Linux there is up2date, yum and apt-get. Sun provides and ftp site from which to download patches.

Windows

For Windows, there are several solutions. First, Microsoft provides an “Automatic update” service that can be triggered manually through Internet Explorer, or via the control panel it can be configured to update automatically.

There are other ways to update a Windows computer. Microsoft provides two types of servers that can provide updates. The first is called SMS, or Systems Management Server -- www.microsoft.com/smsserver/evaluation/overview/default.asp. This works in cooperation with an Active Directory to control the updating of systems on the domain. The second is called SUS, Software Update Service, www.microsoft.com/windowsserversystem/sus/default.msp, which uses a web based technology for clients to check into the SUS server for possible patches or updates.

Linux

There are several update technologies for Linux. There is RedHat’s up2date, Fedora.us’s YUM, and debian’s “apt-get” service.

up2date

RedHat’s up2date service provides updates and patches for currently supported RedHat products, and the cutting edge Linux distribution Fedora that feeds the “pay for”, stable and supported RedHat distributions. A fee is charged for a subscription to the RedHat Enterprise versions of the update service, although there is no fee for the Fedora distribution update service.

To initialize the up2date service in fedora, boot into run level three, by doing the following as root:

```
# init 3
```

This will go into text mode, which is necessary because some of the updates will likely be for the Xwindows system that can't be updated while running.

At the command line, as root, do the following:

```
# up2date
# rpm --import /usr/share/rhn/RPM_GPG_KEY
# rpm --import /usr/share/rhn/RPM_GPG_KEY_fedora
# up2date -u
```

If you want to update the kernel as well, do the following:

```
# up2date -uf
```

For more information on up2date, type "man up2date".

To go back to xwindows mode, type the following as root:

```
# init 5
```

yum

Yum is included in Fedora's Linux distribution. This project is based on Drexel University's linux.duke.edu/projects/yum/ modified yellow dog updater. To update RedHat based systems, including versions 8, 9 Fedora 1 and 2, a yum rpm package can be downloaded from www.fedora.us/wiki/FedoraHOWTO to connect to Fedora.us's update service at the University of Hawi Information and Computer Sciences department.

To start the update process go to run level 3. Xwindows needs to be updated as well as the rest of the system and can not be running while being updated.

```
init 3
```

then start the update by doing the following:

```
yum update
```

Yum, as installed by default, works without further configuration, unlike up2date. Yum will download the most recent rpm headers and compare them to the previously installed rpm headers, and if they are newer, it downloads and installs the newer package. Yum also checks for dependencies. For instance, ssh depends on zlib and openssl packages, and those two files must be at a certain revision level for the new ssh to work properly. So yum will get and install the other two packages before installing the new ssh.

apt-get

The apt-get, or Advanced Packaging Tool, was created by the Debian Linux developers. Since then, a port has been made to manage rpms as well as Debian's deb packages. A link to a "HowTo" use APT can be found at: www.debian.org/doc/manuals/apt-howto/index.en.html. To use it for rpm packages on RedHat 8, 9, Fedora Core 1 and Fedora Core 2, go to the University of Hawaii's www.fedora.us/wiki/FedoraHOWTO.

Solaris

One longstanding method to update a Solaris machine is to use the ftp site that provides current recommended patches. The steps to deploy patches on a Solaris system are as follows.

Use the "ftp -p sunsolve.sun.com", log in anonymously and "cd /pub/patches". Set common download parameters such as "bin", "hash" and "prompt". Download the patches via mget - "mget 9_Recommended*" (for Solaris 9) currently about ~100megs. Look through the readme file for patches that might create problems with the system. Unzip and untar the files, then go into the 9_Recommended directory and run "install_cluster". This installs the patches one at a time, using the pkgadd program. This can take from twenty minutes to over two hours depending on how fast the machine is. The two most common errors are: "error 2: patch already applied" and "error 8- patch not needed by system". Do not forget to reboot the system for the patches to take effect.

There are two useful things to know after patching a Solaris system. First, the "showrev -p" command will show all the patches that have been applied to the system. The second is that logs and information about the patches applied are kept in /var/sadm/install_data/*.logs.

To find out information on patches for Solaris, Sun provides the following website: sunsolve.sun.com/pub-cgi/show.pl?target=patches/patch-access.

In Solaris 9, it is possible to use a combination of the GUI "smc" command and the Solaris PatchPro to automate the patching process. Solaris less than 9, may use the PatchPro package can also automate the patching process, however without such a nice interface.

Sun has a java based PatchPro expert web page, however, be cautioned to know the computer system fairly intimately before attempting this. patchpro.sun.com/servlet/com.sun.patchpro.servlet.StorageServlet?referrer=classic.sunsolve.sun.com&form=expert/expert.html

Special cases

There are several computer related devices that also need to be patched and secured. Personal routers, networked printers, home DSL or Cable routers, Wireless access points also require updates. These devices are generally vulnerable out of the box and require configuration to be secure. When vendors catch problems with their hardware they generally provide firmware updates and instructions on how to apply them. Often they not only fix security problems, but

also performance issues with the device. To find firmware updates for a device, go to the vendors' website and search for the model of the device you have, and there should be a link to download firmware updates.

An example, for the dlink DCM-202 cable modem, dlink provides the following firmware upgrade: support.dlink.com/products/view.asp?productid=DCM%2D202

© SANS Institute 2005, Author retains full rights.

Data Security

How is your data protected? Is regular access control enough to protect the data? Often it is not. There are programs that will hide or encrypt data so it is not readable without the appropriate software or “key” to unhide or unlock it.

Stegonography

Even though stegonography is not actually encryption, it does fall into that category. Stegonography is the art of hiding things in plain sight. Stegonography makes an attempt to use the noise in a system, such as a picture, and replace that noise with data.

Encryption

Encryption is the art of creating a code to change a message so it is unreadable, then using the code to make the message readable again. This could just be shifting the alphabet over three letters, then using this shift to create the message, then the person decrypting the message just uses that shift to decrypt the message. This method of rotating the alphabet is called a “rot-3”, and was used by Julius Caesar. It is also known as the Caesar Cipher, ITsecurity website has a good definition and example of this at: www.itsecurity.com/dictionary/caesar.htm¹⁵.

Another explanation of encryption is that it is a method to take the data, and with a special key, run it through a mathematical algorithm to change it so it is not recognizable. A key is required to unlock the data so it is again in a useful format. There are two basic types of encryption, Synchronous and Asynchronous.

Synchronous encryption is when the key to lock, and the key to unlock the data are both the same. Asynchronous encryption is when the key to lock the data, otherwise known as a private key is different from the key to unlock the data, which is known as a public key.

Modern encryption techniques use a combination of both synchronous and asynchronous techniques to make a more complicated and harder to break encryption.

Resources

Some resources on the web on the subject of stegonography and encryption are:

www.microsoft.com/windowsxp/using/security/learnmore/encryptdata.mspx
www.invisiblesecrets.com/invsecre4.html www.gnupg.org/ www.pgp.com
netsecurity.about.com/cs/hackertools/a/aafreecrypt.htm www.rsasecurity.com
www.cypherix.com/cryptainerle netsecurity.about.com/cs/hackertools/a/aafreecrypt.htm
mrcorp.infosecwriters.com/Steganography.htm www.netip.com/links/stegonography.htm
www.thefreecountry.com/sourcecode/encryption.shtml

¹⁵ www.itsecurity.com/dictionary/caesar.htm, Copyright (c) 1999-2004 ITsecurity.com, All Rights Reserved. Devon, UK, General Editor: Kevin Townsend, Dictionary Editor: Dave Shore

Backups

If the computer system has been compromised, and data integrity on the machine is destroyed, how do we recover?

How important is your data? Would you be upset at losing years worth of e-mail, documents and/or research? Are you storing electronic data that has corporate or technological trade secrets? Are you storing personnel data? Do you use or store any other kind of important or secret data? Are there a lot of applications installed on your computer? How long did it take to set up your computer? These are crucial questions to ask when deciding on a backup strategy.

There are various technologies that are available for creation of backup strategies. Each has its advantages and disadvantages.

Software

There are several software technologies that can be used to back up a computer.

Norton Symantec Ghost

This product can back up NTFS, FAT (Windows file systems) as well as Linux file systems. It can create a complete disk or partition images of the drive, so it can be restored exactly to the state of when the backup was made. When restoring Linux images, learn about how to restore the boot sector using lilo or grub. More information on Ghost can be found at:

www.symantec.com/sabu/ghost/ghost_personal

Unix backups

There are a couple of ways to create an image of a hard drive in *nix related systems. First is the “tar” command. This stands for Tape ARchive. The second is the combination of the dump and restore commands. These can provide different levels of backup. A level 0 backup does a complete, full backup. A level 8 backup backs up everything that has changed since the last level 8 or lower backup has been done. A level 9 backup will back up everything that has changed since the last level 9 or lower backup has been done.

There are more software technologies that provide backup services, one of which is Retrospect, which can be found at www.dantz.com.

Remote Backup services

There are companies that provide offsite backup services for a fee. A few such companies are:

www.livevault.com

www.usdatatrust.com/remote_backup.asp

oneeach.com/remote_backup.html

Hardware

There are several hardware technologies that are available for backing up data.

USB pen drives

These devices come in sizes from 16 megabytes to 4 gigabytes currently.

Advantages:

The convenience of this device makes it easy to transfer data from one computer to another, and the media is small and easily placed in a pocket or hung around your neck. Some of these devices are bootable, which is useful in case the hard drive has critical OS files either corrupted or missing. If you need to leave the area because of a fire or natural disaster, the data is easy to take with you.

Disadvantages:

If a usb pen drive is physically removed from the computer, and not "ejected" or "umounted" then often the files that were copied or saved onto the device will not be found. What happens is the operating system will put the files in a buffer, or temporary holding place in memory, and when the request is made to the operating system to "eject" or "umount" the drive, the OS flushes the buffer to the usb pen drive and completes the write. Another problem with usb pen drives is their rate of failure. They are fragile devices and it is easy to damage them. If the drive is not properly ejected or umounted, this will also increase the rate of failure. Even if properly used, they don't have a long life span. Another disadvantage is that it is easy to conceal a usb pen drive and people visiting your organization, or even discontented employees, can come in the building with a 4Gb pen drive and leave with sensitive data that could cost significant amounts of money and woe if lost. Lastly, USB drives do not hold as much information as can be held on a hard drive or other large media.

For more information and statistics on USB pen drives, see the arstechnica.com/reviews/hardware/flash.ars/1 review on the web.

CD/DVDs

CD/DVD recordable drives can be found for a wide range of prices and technologies.

Advantages:

CD and DVD media is much cheaper than USB pen drives. If treated well, and a good brand is chosen, these will last for a long time. Every computer now days, has a CD or DVD player, so moving data from one computer to another is easy. This media can be used as "read only" or "read-write". In read-write mode a CD can be written to hundreds of times. CD/DVD's are also easy to transport in case of emergency.

Disadvantages:

This media is easy to damage as well and must be kept in a case or risk being scratched. A scratch can render the data, whether a movie, music, presentations, documents or other data, unusable. Just as with USB drives, CD/DVD's do not hold large amounts of data that can be held on hard drives and other large media.

Hard drives

Hard drives can also be used as a backup medium. A direct mirror can be made of the original hard drive by copying the data off it to onto backup hard drive. The price depends on the speed and technology that the drive uses.

Maxtor has a "one-touch" technology available to use a button on an external hard drive to back up the computer. This comes in firewire and USB technology.

Advantages of Mirroring

If the primary hard drive fails, very little needs to be done to replace it with the backup and be running again and the rate of failure is much less than the USB drive or CD/DVD's. The primary reason for this is, once the hard drive is installed, it is not moved, removed or carried around like the other devices. A real advantage of this is that everything that is held on the primary drive can also fit on the backup drive. It is the fastest method to be up and running after a crash.

Disadvantages of Mirroring

If there is a catastrophic disaster and a fire or other natural disaster destroys the computer, both primary system and backup are destroyed.

Advantages of One Touch External Hard Drives

This type of external drive uses software technology similar to that which is used for tape backups. Backup software creates an original full backup, then for a week or so, only backs up files that change, then make a full backup again the next week. Complex schedules with different levels of backups can be made. This, like USB flash drives and CD/DVD's can easily be taken out of the building when there is an emergency.

Disadvantages of One Touch External Hard Drives

Software needs to be installed to take advantage of the “one-button” technology. The backups can not contain a “live” OS, just backups of files. If the computer fails and files need to be restored, first an operating system must be installed, then the backup software installed before files can be restored.

Tapes

Tape drives come in a variety of sizes and technologies. Some drives handle a single tape, some “tape libraries” can handle hundreds of tapes. The cost depends on the technology, and how many tapes will be used. In most cases, a single tape can take care of a single computer.

An article on “Tom’s Hardware” website describes and compares different tape technology, and it can be found at: www.tomshardware.com/storage/20040408/tandberg-streamer-01.html

Advantages:

Tapes can hold a large amount of data, and are easily transported in case of emergency.

Disadvantages:

Transportable media can be dropped or broken. As with the “one-touch” hard drive, an operating system would need to be installed before data could be restored from tape. Initially, they are an expensive backup option.

© SANS Institute 2006. Author retains full rights.

Disaster Recovery Plan

Disaster recovery is often not taken seriously, or not thought of at all. How important is your data? If a disaster happened, how long could you go without your backed up data? Is it important to have an offsite backup? A plan needs to be written and recorded so that if something does happen, the computer or computer systems can be restored to a reliable state as close as possible to pre disaster conditions quickly.

First, determine possible points of failure. How reliable is the power grid you are connected to? How reliable is the network you are connected to? Do you rely on any networked servers or on any applications that are served over the network? What would happen if your building was destroyed? Does your data reside on networked disks?

Classify the importance of the data. The operating system, for instance, can easily be replaced. Special servers, services, customer, proprietary or personal information all have value. What would happen if it were to permanently disappear? How quickly could you retrieve it? How much space does the information you need backed up take, and what media will it fit on?

Determine what steps can be taken to mitigate possible points of failure. Can the room or building the computer(s) are in be flood-proofed? What happens if there is a fire in the room where your computer is? What happens if there is a fire in the area that might consume the building? What can be done if the power to the computer is unreliable, which could blow a power supply?

A plan is needed for the recovery of the data. Will new computer hardware need to be acquired? Is a written disaster recovery plan in place, which is available off-site in case an emergency or disaster happens?

A plan is needed for the resumption of regular activities, once the data has been recovered. How will data be validated? Have checklists or written procedures been followed to recover data? Have the correct versions of software been installed so that what worked before will work again?

Assume the worst happens, what can be done to prepare for it? Do you need off-site backups? Do you need a separate site that can be turned on and be ready to go as soon as the emergency or disaster takes place?

More in-depth strategies and papers on HowTo's for Disaster Recovery plans go about putting together a plan can be found on the web at: www.sans.org/rr under the Disaster Recovery category.

Keep Up With Computer Security Technology and News

There are several ways to keep up with computer security technology. There are many resources on the web that contain current computer security news, white papers, HowTo's and vulnerability databases.

One computer security news site is www.securityfocus.com. This site covers current computer security news, opinions from many specialists in the computers security arena, information on vulnerabilities, "how to" articles explaining current and new technologies, and books on computer security topics. Current information on computer security news can be found at isc.sans.org. This site gives daily information on computer vulnerabilities, internet problems, product reviews and security suggestions. Another site that contains computer security news is www.securiteam.com, which hosts news on a variety of computer security products, vulnerabilities and exploits.

There are sites that cover white papers and howto's. One site that hosts research papers written by students seeking computer security certification is www.sans.org/rr. This is the "reading room" of the SANS organization that provides training on many computer security issues. This site covers many computer operating systems, technologies and networking topics. www.securitydocs.com houses a variety of documentation and information on computer security.

For Linux operating systems, and closely applicable to other *nix operating systems, "The Linux Documentation Project", www.tldp.org, houses a vast assortment of HowTo's covering many topics on installing, configuring and maintaining a computer system. For Windows operating systems, Microsoft has some online documentation as well at: support.microsoft.com/default.aspx?scid=fh;EN-US;FAQS.

There are several online resources that give information on vulnerabilities and exploits. secunia.com hosts advisories on latest security vulnerabilities of many products, and news on viruses. The Open Source Vulnerability Database, www.osvdb.org gives information on vulnerabilities that can be found in open source software that is used in a wide variety of operating systems and platforms. Organized by the CIO of the DOE, the CIAC, Computer Incident and Advisory Capability - www.ciac.org/ciac/index.html, provides posts of technical bulletins and vulnerabilities on the above website. US-CERT, www.us-cert.gov/, stands for US-Computer Emergency Readiness Team, is cooperation between government and the private sector established to protect US computer infrastructure.

Host Based Cyber Defense Strategy Review

Do you connect to the internet? Do you conduct any business online? How important is the data that is on your computer? Choices must be made to determine the level of acceptable risk that can be tolerated on your computer system or network.

Controlling network traffic with scanners, filters, intrusion detection/prevention devices is important. However, without host based security, any computer security plan is incomplete.

A viable plan must include multiple layers and angles of defense. Minimizing system services, controlling access, logging, scanning for malware, patching, securing data, backing up data and maintaining a disaster recovery plan are essential for everyone – from the home user to the largest corporation. The computer security field is developing and expanding and to minimize vulnerability, it is vital to keep up with technology and news.

© SANS Institute 2005, Author retains full rights.

Bibliography

- (1) Frisch, Aileen, Essential System Administration, Second Edition. O'Reilly www.oreilly.com, 1995: 592 – 595, www.oreilly.com/catalog/esa2/index.html
- (2) Sun Microsystems, System Administration Guide: Basic Administration. docs.sun.com/db/doc/817-6958/6mmafc30c?a=view
- (3) Wirzenius, Lars; Oja, Joanna et al, The Linux System Administrator's Guide Version 0.8. The Linux Documentation Project, 2003: 9.3 at the following website: www.tldp.org/LDP/sag/html/x2140.html
- (4) Frisch, Aileen, Essential System Administration, Second Edition. O'Reilly www.oreilly.com, 1995: 156, www.oreilly.com/catalog/esa2/index.html
- (5) Man page for login(1) on the Sun Product Documentation website: docs.sun.com/db/doc/816-0210/6m6nb7mdo?a=view
- (6) Man page for login(5) on the Linux Valley mirror of The Linux Documentation Project website: www.linuxvalley.it/encyclopedia/ldp/manpage/man5/login.defs.5.php
- (7) Frisch, Aileen, Essential System Administration, Second Edition. O'Reilly www.oreilly.com, 1995: 217, www.oreilly.com/catalog/esa2/index.html
- (8) Bradford, Ed and Mauget, Lou, Linux and Windows Interoperability Guide. Prentice Hall PTR www.phptr.com, 2002: 365 - 369, www.phptr.com/title/0130324779
- (9) Kernighan, Brian and Pike, Rob, The UNIX Programming Environment. Prentice-Hall, 1984: 53, www.amazon.com/exec/obidos/tg/detail/-/013937681X/102-2335063-1351366?v=glance
- (10) Kernighan, Brian and Pike, Rob, The UNIX Programming Environment. Prentice-Hall, 1984: 55, www.amazon.com/exec/obidos/tg/detail/-/013937681X/102-2335063-1351366?v=glance
- (11) Barrett, Daniel J., Silverman, Richard E. and Byrnes, Robert G., Linux Security Cookbook. O'Reilly www.oreilly.com, 2003: 23 - 48, www.oreilly.com/catalog/linuxsckbk/index.html
- (12) Reuters, “\$55bn virus damage costs for businesses last year”. www.silicon.com, 2004, www.silicon.com/software/security/0,39024655,39117842,00.htm
- (13) Bridis, Ted “Group, Dell Launch Anti-Spyware Campaign”, Associated Press, story.news.yahoo.com/news?tmpl=story&u=/ap/20041016/ap_on_hi_te/internet_spyware&e=1
- (14) Zenkin, Denis “Understanding Macro Viruses”. www.securityfocus.com, 2001, www.securityfocus.com/infocus/1278
- (15) www.itsecurity.com/dictionary/caesar.htm, Copyright (c) 1999-2004 ITsecurity.com, All Rights Reserved. Devon, UK, General Editor: Kevin Townsend, Dictionary Editor: Dave Shore

Web sites

Anti-malware

www.kaspersky.com
www.lavasoftusa.com
www.rkhunter.org

www.symantec.com
www.spybot.info/en/index.html

www.mcafee.com
www.chkrootkit.org

Backups

www.symantec.com/sabu/ghost/ghost_personal
www.usdatatrust.com/remote_backup.asp
oneeach.com/remote_backup.html

www.dantz.com

www.livevault.com

Computer Documentation, HowTos and Whitepapers

www.tldp.org
www.sans.org/rr
www.theeldergeek.com
www.securitydocs.com
www.obfuscation.org/ipf

www.microsoft.com/windowsxp/support/default.msp
www.sun.com/documentation
www.linuxforum.com
www.computerproblems.org
support.microsoft.com/default.aspx?scid=fh;EN-US;FAQS

Encryption

www.InvisibleSecrets.com
www.gnupg.org
www.rsasecurity.com
www.pgp.com
www.netip.com

netsecurity.about.com/cs/hackertools/a/aafreecrypt.htm
mrcorp.infosecwriters.com/Steganography.htm
www.thefreecountry.com/sourcecode/encryption.shtml
www.cypherix.com

Firewalls

coombs.anu.edu.au/~avalon

www.iptables.org

Internet Encyclopedias and dictionaries

wi-fiplanet.webopedia.com

www.itsecurity.com/dictionary/dictionary.htm

Package Managers

www.rpm.org
www.debian.org/doc/manuals/apt-howto/index.en.html

www.fedora.us/wiki/FedoraHOWTO

Patching

windowsupdate.microsoft.com
linux.duke.edu/projects/yum
patchpro.sun.com

www.sgi.com/support/security/patches.html
www.fedora.us/wiki/FedoraHOWTO
www.debian.org/doc/manuals/apt-howto/index.en.html

Security News

www.linuxsecurity.com

www.securityfocus.com

www.us-cert.gov

secunia.com

www.securityfocus.com

www.securiteam.com

www.virusbtn.com/news/latest_news/index.xml

www.ciac.org/ciac/index.html

www.cert.org

netsecurity.about.com

isc.sans.org

www.osvdb.org

SSH Applications

www.openssh.com

www.f-secure.com/products

Operating Systems

www.microsoft.com

www.redhat.com

www.sun.com

www.debian.org

www.slackware.org

www.adminschoice.com

Shopping for Computer Supplies

www.pricegrabber.com

www.pcconnection.com

www.mwave.com

www.techforless.com

www.buy.com

www.zoomtek.com

www.provantage.com

Technology News and Reviews

www.arstechnica.com

www.tomshardware.com

www.slashdot.org

Trace File Decoders

www.ethereal.com

Tripwires

www.tripwire.com

www.tripwire.org

sourceforge.net/projects/aide