



© 2003. All rights reserved.
Black Box Corporation.

BLACK BOX[®]

NETWORK SERVICES

Network Security

A white paper

1.0 Introduction

In today's world it's impractical—if not impossible—to isolate your network. Networks are interconnected and usually feature a full-time connection to the Internet, which can be a dangerous place. Although the vast majority of Internet users are honest and benign, there are a few who, motivated by greed or maliciousness, will attempt to directly or indirectly hurt you or your network. They can destroy your operating system, damage data, steal information, overwhelm Web sites, and clog e-mail servers. Even if your network isn't connected to the Internet or other networks, it's still vulnerable to attacks transported by diskettes or CDs, or from individuals with access to your network.

Attacks are surprisingly common. Virus or worm outbreaks large enough to make the news happen nearly every week and many, many more go unreported except at the Web sites of anti-virus software vendors. Corporate firewalls routinely log thousands of probes by hackers every day.

This white paper will give you a basic understanding of security threats and give you an overview of how to defend your network. It will not make specific recommendations as to what your security policy should be. Every network in every organization has different security needs, and security strategies must be adapted to each specific situation. Because it's not always easy to distinguish between threats to individual users and threats to networks, this paper will cover the whole range of threats from simple scams to large-scale, denial-of-service attacks.

2.0 Threats to your network

2.1 Viruses

Viruses are bits of software code—usually disguised as something innocent—that replicate themselves in PCs. The term virus is generally used as a generic word for a wide range of malicious codes, some of which are not technically viruses.

Creating or distributing a virus is classified as a computer crime and is a prosecutable offense. Laws enacted against the creation of viruses include the Electronic Privacy Act of 1986 in the United States and the Computer Misuse Act of 1991 in Europe.

Some viruses need help to reproduce, some replicate all by themselves, but all reproduce without your permission or knowledge. Viruses often have an infection phase where they reproduce and an attack phase where they do damage. The ability to infect varies from virus to virus, as does the damage they can cause.

Viruses are the number one method of computer vandalism with new ones being identified every day. There are more than 70,000 known computer viruses. The damage they cause varies widely and can include no damage at all, a message left on your screen, corrupt data, or complete system shutdowns.

Viruses are roughly categorized into subgroups such as common file-infesting viruses, worms, Trojan horses, macro viruses, and others. There is also an entire family of related hoaxes that may be damaging but aren't actually viruses. These hoaxes usually arrive in the form of an e-mail that warns you against a nonexistent virus or attempts to convince you to do something that will damage your computer.

2.1.1 File-infesting viruses

File-infesting viruses are the most common ones in the virus family. They infect executable files by adding their own code to that of the original file. You get file-infesting viruses on your computer when you run an executable file—a game, for instance—received in your e-mail or brought in on hard media such as a floppy disk. As soon as you run the infected file, the virus attaches itself to other executable files on your hard drive. When you transfer infected files to another computer, the virus goes along for the ride and finds more files to infect. A variation of the file-infesting virus infects a file such as an .OVL or .DDL file that is called by an executable file but doesn't infect the executable file itself.

Because they require the transfer of an executable file, which most people know better than to open, common file-infesting viruses tend to spread relatively slowly and don't cause the wildfire infections across the Internet that worms do.

2.1.2 Worms

Worms are self-replicating viruses. Unlike an ordinary virus, which depends on the transfer of a host file in order to replicate, a worm is an independent entity that usually spreads itself without needing a computer user to transfer a file. Many of the newer computer viruses that make the news are worms. Most travel primarily by e-mail but some also spread through nontraditional means such as the Internet Relay Chat (IRC), peer-to-peer networks, or even Web sites—it's possible to be infected by a worm simply by visiting a compromised Web site.

Worms tend to spread very rapidly and can cause a lot of damage—an August 2003 world-wide attack of the Blaster worm shut down corporate networks, slowed air traffic, and even took down the U.S. Navy Marine Corps Intranet (NMCI). In September 2003, the Swen worm infected 1.5 million computers and clogged e-mail boxes around the world with fake Microsoft patches, causing massive inconvenience even for uninfected users.

2.1.3 Trojan horses

A Trojan horse does not replicate itself and is technically not a virus. A Trojan horse (also known as a Trojan) is a software program, often a game or utility, that seems to do one thing but has incorporated within itself another, secret function that will cause damage, pass on information about your computer, or enable its sender to hijack your computer.

Trojans are often part of hybrid or multipartite viruses. For instance a Trojan may be “planted” in an application by a worm or it may incorporate a virus within itself.

Sometimes a Trojan horse will cause obvious damage, but often it looks and acts like a legitimate program, hiding its illicit work in the background where you don't notice it.

Some Trojan horses are used to spy on your activities and steal information, however, the hackers that use Trojan horses to get to your PC are usually not at all interested in your data. What they're after is a PC with an “always up” broadband connection to the Internet. When they find a vulnerable PC, they use it to distribute pornography or spam e-mail or to launch a Distributed Denial of Service (DDoS) attack in which they use remote PCs to attack a Web site or e-mail server. The reason they use your PC as a host for these activities is because, to the recipient, it looks like the unwanted traffic is coming from your machine. This hides the culprits from ISPs and law enforcement, making them nearly impossible to find.

Trojan horses can be very sneaky—if you're not looking, you may never notice that your computer has been used this way.

A common Trojan horse is the Backdoor Trojan which comes in many variations, usually enabling a hacker to access your computer through the IRC.

2.1.4 Macro viruses

Macro viruses are written in the internal macro language provided with many applications. A macro is a set of instructions within a program that you record and assign to a short key code. Then, when you type the key code later, the recorded instructions execute. Macros can simplify day-to-day operations such as formatting documents because you can record and play back repetitive instructions instead of having to do them yourself one by one.

Many programs enable you to extend the macro language provided with an application with more complex programming languages, so you can set up very complex routines as macros. It's not only possible, it's downright easy to write viruses or modify existing viruses within this format. Macro viruses are extremely common—especially within Microsoft Excel and Microsoft Word files—with thousands of variants identified. They spread easily because they travel in documents, which are often shared—unlike executable files—and also because many of them, in wormlike fashion, e-mail themselves to everyone they find in your address book. Because macro viruses are so easy to write and modify, new ones pop up all the time. Additionally, because macro viruses spread within an application, they may spread between operating systems—for instance from a PC to a Mac.

Macro viruses activate when an infected document is opened. The degree of damage varies depending on the particular virus. Many macro viruses are relatively harmless—one merely inserts the word “Wazzu” into text—but some can do serious damage by deleting files or formatting disks. A recent macro virus is the Melissa virus, which propagates itself by mailing itself in the form of an e-mail message usually called “Important Message From <your name>” containing an infected Word document as an attachment. Usually the message in the e-mail is “Here is that document you asked for ... don't show anyone else ;-)” Unsuspecting victims think they're getting a file from a friend and open the document, which then infects any templates it finds and mails itself to everyone in their address book, and so on.

2.1.5 *Boot sector viruses*

Boot sector viruses infect the boot sector on a floppy disk or the Master Boot Record (MBR) on a hard drive by overwriting the original boot code with its own code. The boot sector or MBR generally resides on the first sectors of your hard disk and controls the boot sequence when you start up your computer. A virus that infects these sectors is especially dangerous because every time you start up your computer, it's loaded into memory from where it can spread to other parts of the hard disk or to another disk. Boot sector viruses frequently cause a complete system failure in which your PC can't start up or find its hard drive.

2.1.6 *Stealth viruses*

Stealth viruses go to great lengths to actively hide themselves by residing in memory while running. They fool the operating system by modifying and forging the results of calls to functions in the infected file, so the system believes it's reading the original file. A stealth virus can even hide the fact that it's consuming memory. Stealth viruses hide themselves so well, they can sometimes fool antivirus products into thinking a computer is virus free.

The Brain virus, discovered and catapulted to brief fame in 1986, was the earliest known example of a stealth virus. This virus infects the boot sectors of a floppy diskette, redirecting the operating system when it attempts to read the boot sector.

2.1.7 *Polymorphic viruses*

A polymorphic virus attempts to defeat virus-scanning software by using an algorithm to encrypt itself each time it infects a new host. The encrypted virus escapes detection by the anti-virus software and then decrypts itself to infect the computer. This virus is very difficult to detect because its signature is different every time it infects a new host. More sophisticated polymorphic viruses vary their encryption methods, making them even more difficult to detect.

2.1.8 *Multipartite viruses*

Multipartite viruses are the hybrids of the virus world—they're most commonly a combination of boot sector viruses and file infecting viruses, and infect both system sectors and files. These viruses are fairly rare because they're difficult to write but tend to be particularly nasty when they occur. The famous Hare virus of 1996 (famous for being the most publicized virus that never spread) is an example of a polymorphic multipartite virus.

2.2 *Invasive software*

Adware, spyware, scumware, drug dealer ware, and theftware are all kinds of software that arrive on your computer—sometimes with your permission—and set about the business of polluting your screen with ads, sending information about you out onto the Internet, and generally making a nuisance of themselves by slowing down your system or even causing system crashes.

Some of this nasty software is actually more like a stealth virus or Trojan horse, loading itself without your permission, hiding out in your system, and resisting removal. However, much of it you agree to and download yourself without really knowing what it is you're downloading.

2.2.1 *Adware*

Adware is commonly acquired when you download it yourself from the Internet. Freeware or shareware—software available for free or for a small fee on the internet—often comes with hidden adware. Adware has even turned up in software sold through normal commercial channels.

In the usual scenario, you decide to download a program—often a game or useful utility—that's available free on the Internet. When you download the software, you have to click through a page of unintelligible legal stuff that includes things like copyright and, in the very fine print, permission to install adware along with the software. You click okay to download your software. Soon you start noticing that you're getting more popup windows than usual. What you may not notice is that your PC is reporting your Web browsing habits to someone out on the Internet.

Some adware is produced by legitimate companies that are quite open and honest about what they're providing. For instance, Gator, a utility that helps you fill out forms, tells you that it includes adware, makes you click through multiple clearly written disclosure screens before you install it, tells you exactly what it does, has an icon on-screen when it's in use, and provides instructions for removal.

However, a lot of adware comes from distributors who squirrel it away inside totally unrelated software and get your "permission" with a page of virtually unreadable legalese. The popular music-download software, Kazaa, was in the news for incorporating software that could hijack your computer to be a part of a large "super network." This rather alarming feature was mentioned in the fine print of the disclosure form, but almost none of Kazaa's users read and understood the form.

The fact is, clear, readable terms of service and privacy contracts are rare. Every day, thousands of people agree to terms they haven't read or don't understand, and they download software they never intended to.

Adware ranges from helpful to unethical, but it's still much better than spyware.

2.2.2 *Spyware*

Spyware sneaks into your computer without bothering to ask for your permission. As with adware, you may download it along with other software, except spyware doesn't even have a fine-print page that mentions you're also getting unwanted features. Some spyware even does a "drive-by download" and installs itself when you just visit a Web page—sometimes as a mere cookie, sometimes even as a complete application. Much spyware is more properly classified as a worm or Trojan horse.

Spyware often hides itself so that it's difficult to find and eradicate from your hard drive. Some spyware will resist being removed and even take parts of your operating system with it when you attempt to eradicate it.

Spyware usually has many of the same functions as adware, putting up ads while you're browsing the Web. However, spyware, is far more likely than adware to engage in antisocial behavior such as sending personal information to a third party, resetting your home page or altering system files. One of the nastier features that can show up in spyware is key logging, which enables it to record anything you type, including your passwords, e-mail messages, real-time chats, and credit card numbers. Some spyware can even spy on you using your own Webcam.

2.2.3 *Other unpleasant software options*

Adware and spyware have some unsavory relatives. Scumware is designed to steal revenue and traffic from legitimate Web sites. Drug dealer ware offers free software and then shuts down and demands payment months later when you've presumably gotten used to using it. Theftware hijacks ad space on Web pages, replacing the page's ads with its own ads. Malware is an outright virus and is designed to damage or disrupt your system.

2.3 *Spam*

If you have e-mail, you know that spam isn't the same as Spam[®], a delicious, pork-based, canned meat product sold by Hormel Corporation. Spam on the Internet is the junk e-mail that arrives in floods and torrents. It clogs servers and in-boxes with offers for dubious products, links to pornographic sites, chain letters, hoaxes, and scams of all kinds. At best, spam is merely a nuisance, but if you're not paying attention, spam can also be dangerous.

Spammers usually get their address lists by mining usenet newsgroups and searching the Web for e-mail addresses. Sometimes they will buy an address list from another spammer.

Because most ISPs block e-mail coming from known spammers, many spammers route their mail through unsuspecting third-party servers. This tactic jams the intermediary's e-mail server with unwanted spam messages and, worse, makes them a target for complaints about the spam that seemed to originate with them but didn't.

Another common tactic spammers use is to open a "disposable" trial Internet account at an ISP. By the time the ISP realizes they've been had, the spammer has already sent tens of thousands of e-mails. The unsuspecting ISP is left to clean up the resulting mess.

2.3.1 Advertising spam

Most spam is advertising. And most of what spam advertises is deceptive or fraudulent. Popular advertisements include get-rich-quick business opportunities, work-at-home schemes, miracle cures, dubious investments, illegal cable descrambler kits, credit and credit repair offers, vacation prize scams, and sexually explicit ads.

Spammers send advertising through e-mail because it's cheap. When you're marketing what's undesirable, it helps to advertise in a place where you don't have to pay for the ads. The cost of spamming is so low, there's no motivation to target the ads to specific customers. Why bother to do any market research when you can hit a million e-mail addresses for the same price?

2.3.2 Dangerous spam

Some spam will bite you by carrying viruses or by trying to lure you into providing credit card numbers.

Some dangerous "spam" e-mail comes, not from spammers, but from worms, that generate infected e-mail from the computers of unsuspecting hosts. For instance, the recent invasive Swen infestation clogged e-mail boxes to the point that many people found in-boxes stuffed with more e-mails that were infected than not.

Scams that try to part you from your money are also quite popular on the Internet. Recently many people got e-mails that appeared to be from eBay or PayPal telling them that they needed to update their accounts. The e-mails looked real, complete with the appropriate logos, and provided a link to what looked like legitimate eBay or PayPal Web pages. Unsuspecting users who filled in the requested information found their credit card numbers and sometimes their very identities stolen. Con men and thieves have entered the high-tech age and are using Internet spam as a cheap way to run their con games.

2.3.3 Hoax spam

Some spam is propagated, not by professional spammers or viruses, but by well-meaning people just like you who pass on chain letters, pleas for the return of "missing" children, and fake virus alerts.

Chain letters are letters that promise a large return for small effort and often also threaten bad luck if you break the chain. If they involve the exchange of money, they're not only a nuisance, they're illegal.

Heart-tugging appeals to our sense of altruism also account for a lot of spam e-mail. Often these urge us to look for missing children or send e-mails to sick children. A variation appeals to our sense of outrage. For instance, one claims that the government is going to tax e-mails; another claims that PETA is dressing deer in orange hunting vests.

There are malicious hoax e-mails out there, too. Some hoax e-mail tries to get you to do something to harm your system under the guise of "helping" you. Fake virus warnings warn you about viruses that don't exist, often while carrying a virus themselves. In an act of ultimate chutzpah, a virus-infected e-mail will sometimes even offer to sell you virus-blocking software to eliminate the virus that it just infected you with.

One sign that an e-mail is a hoax is that they almost always urge their readers to forward them on to "everyone in your e-mail list." Something for nothing—forward this e-mail and Microsoft will send you money—is also a common theme of this sort of hoax.

2.4 Hackers

At the beginning of the computer age, way back in the Jurassic...er...the 80s, the word “hacker” referred to a person who was very good with computers. Today, the same word usually refers to a person who maliciously breaks into networks, breaks the security on application software, or creates viruses.

The hackers we hear the most about are the system crackers who hack into individual computers or into networks. They’ve broken into networks at banks, universities, and military bases. They’ve infiltrated networks at the Pentagon. Corporate networks are common victims—most have been hit by hackers at least once.

Hackers have changed Web sites, accessed unauthorized services, altered and removed files, and stolen information. Often a hacker breaks in just to see if he can do it, looks around, and leaves without doing any damage. Sometimes the damage is minimal—a common trick is to alter the victim’s Web site. However, hackers can and do cause very serious damage to systems and have even been known to erase evidence of their activities by altering logs, leaving network administrators unaware they’ve been hacked.

Hackers often use known security holes in operating systems—particularly Windows—to penetrate your network. These holes are well-known and are repairable with software patches. For instance, Windows security holes are being discovered almost weekly, and the patches may be downloaded at the Microsoft Web site.

Hackers find security holes to exploit by pinging your network to find an open IP address and then probing to find an open port. Many corporate firewalls are probed thousands of times a day by hackers looking for a security hole.

Some network attacks don’t require that a hacker actually break into a network. Denial-of-service attacks swamp your Web server with more requests for connections than it can handle, causing it to become so slow as to be unusable or even crash entirely. Mail bombs do the same thing to e-mail servers by bombarding them with more e-mail than they can handle.

3.0 Building effective network security

“It is easy to run a secure computer system. You merely have to disconnect all dial-up connections and permit only direct-wired terminals, put the machine and its terminals in a shielded room, and put a guard at the door.” — F. T. Grampp and R. H. Morris, authors of UNIX operating system security, Bell System Technical Journal, 63(8), October 1984.

Although it’s easy to feel overwhelmed by potential threats to your system, disconnecting your network from the Internet and locking it up in a safe room all by itself really isn’t an option in today’s world. Fortunately, it is possible to secure your network and minimize threats. How much network security you need depends on how large the network is, how much it interacts with the Internet, and what your tolerance for risk is.

3.1 Make a plan

For effective network security, you need a plan. Generally, the larger the network you’re administering, the more formalized the plan should be. For a tiny network with a dialup connection, the entire plan can be to keep the virus software updated and not open any suspicious e-mails. A large enterprise network may require a complex, well-choreographed, thoroughly documented plan implemented after a formal risk assessment and analysis of the network.

Your security plan should include:

- User education — Teach network users how to avoid threats and encourage them to act as “eyes” for the network administrator, reporting anything that looks suspicious.
- Access policies — Control physical access to the network through lock and key or password protection.
- Software — This includes the software required to protect your network and the scheduling of regular updates of both antivirus software and patches issued by software vendors.
- Firewall — If a firewall is needed, consider what kind of a firewall is needed, whether and how to install a DeMilitarized Zone (DMZ) (see Section 3.6.5), and schedule regular reviews of firewall policies.
- Backups — data should be backed up on a regular basis.

3.2 Education

The first line of defense against security threats from the Internet is education and common sense. Keep on top of the latest hoaxes and viruses and make sure your network users know about them. There are many well-known sites that specialize in tracking these things for you—see the appendix to this paper for links.

Staying informed when you connect to the Internet is one of your most important safeguards.

3.2.1 Be suspicious

Be suspicious—a suspicious mind is definitely an advantage on the Internet. Learn to be on the lookout for anything that doesn’t look “right” and encourage your network users to do the same.

3.2.2 *Never reply to or forward spam*

Never reply to spam. Some spam tells you that you can opt out of future e-mails by clicking on a link. Don't do it. This is how spammers verify that your e-mail address is a real working address.

Never forward spam. You don't want it—why would your friends want it? That includes chain letters, even a chain letter from a dear, dear friend cursed with a hundred years of bad luck if you don't forward it.

3.2.3 *Keep an eye out for hoaxes and scams*

There is no lost child named Penny Brown. Microsoft does not have an e-mail tracking program and will not pay you to forward e-mail. No one in Nigeria wants to give you 20 million dollars for the use of your bank account. You have not won a lottery in Uruguay requiring you to send tax payments. Your mobile phone will not contract a virus if you answer a call from "unavailable." eBay does not need your credit card information to update your account nor do you need to give your credit card information to cancel laundry charged to your account. There is no such thing as a bonsai kitten.

Hoaxes run rampant across the Internet. Delete any e-mail that promises money, asks for personal information, asks you to forward or respond to something, or tells you something bad will happen if you don't respond. Reputable companies do not e-mail unsolicited software patches, letters asking for credit card information, or links to Web sites that ask for personal information. Even relatively harmless hoax e-mail of the lost-child variety clogs mail servers and is just plain annoying. Do your friends a favor and don't forward it.

3.2.4 *Use caution when opening e-mail attachments*

The most common way to acquire a virus is in an e-mail attachment. Ideally, you should install an antivirus program at your gateway to keep them from ever showing up in your in-box in the first place. However, a virus filter isn't foolproof and sometimes the critters do sneak in. As a rule, don't open any suspicious attachments. You can consider it suspicious if it's from someone you don't know. If it's an unexpected attachment from someone you do know, check with that person before opening it—some viruses mail themselves and look like an e-mail from a familiar person. If you verify that the e-mail attachment is from a person you know and it's an executable file, check to see where they got it before opening—that neat little game could be a Trojan horse. If the attachment is a document created in a program that supports macros, you will, upon opening the document, usually get a dialog window asking you if you want to start the macro. Choose "no" to prevent possible macro viruses from spreading to your computer.

3.2.5 *Use caution when opening e-mail*

Although most viruses that arrive in e-mail travel as attachments, recently worms have appeared that only require that you open an e-mail to infect your computer. The most infamous of these is the BubbleBoy virus, which carries the subject line "BubbleBoy is back!". If you open or even preview this e-mail, it will forward itself to everyone in your Outlook address book. BubbleBoy and his relatives take advantage of a security hole in Microsoft Outlook and can be easily prevented by keeping up-to-date on software patches. However new variants of this virus using new security holes may appear at any time, so make a point of deleting spam without even looking at it.

3.2.6 *Be careful when downloading software*

As a rule, it's not good practice to download software from the Internet, however, there are some very useful utilities out there that are worth downloading. The best way to download useful-looking, but possibly dangerous, software is to first do an Internet search on it to see what kind of a reputation the software has. Carefully read any licensing agreements before clicking on the "agree" button and consider using software that guards against unwanted adware. Then download the software onto a floppy disk or CD and scan it with antivirus software before allowing it onto your hard drive.

3.2.7 *Don't stray into bad neighborhoods*

The Internet has its bad neighborhoods. If you start poking around on sites that offer pornography, gambling, and too-good-to-be true (It just fell off the back of a truck, really...) deals, you're more likely to be exposed to viruses and scams. Stick with sites you trust.

An easy way to wander into a bad neighborhood is to type a Web site name into your browser carelessly or to type in what you think a site might be called. There are parasite Web sites with almost the same name as the "real" Web site. These sites get their traffic by counting on people carelessly entering the wrong Web address. These parasite sites are very likely to offer pornography or infect your computer with a worm. They may offer the same service as the site you were trying to reach. Sometimes they're just annoying and force you to click your way through dozens of popup windows. Parasite Web sites are generally unpleasant places to visit, so be careful entering Web addresses and take the time to look up addresses if you're not sure.

3.3 *Access policies*

Although this paper focuses on network threats that come from outside your network, be aware that security breaches happen "at home," too. Anyone with network access can steal or damage your data or networking devices. No amount of firewall protection is going to save a server if someone steals it. Take the time to look at who has access to what, keep essential network devices under lock and key, and implement password access to sensitive data.

3.4 Software that protects your network

3.4.1 Antivirus software

A crucial component of your security plan is virus software. There are two major types of antivirus software: scanners and checksummers.

Scanners, the most popular variety of antivirus software, scan your hard drive or scan each file in real time as it's accessed. Scanners work by comparing files to known viruses. They're easy to use but must be kept up-to-date with the latest virus information to remain effective.

Because viruses change files, checksummers look for these changes to find signs of infection. They have the advantage of detecting unknown viruses that a scanner can't detect, however they also have trouble distinguishing between legitimate change and a virus infection. Another marked disadvantage of checksummers is that they can only detect infection after it happens—they're useless for virus *prevention*.

Most antivirus software from major vendors is primarily of the scanner type, some adding checksummers for maximum effectiveness.

There is no such thing as ideal antivirus software, and different products have different strengths and weaknesses. For the most effective protection, it's a good idea to use more than one antivirus program.

Antivirus software should be installed at the gateway — where your network meets the Internet, at the server level, and at the desktop. The software at the gateway screens out most infections before they get into your network. Regular scans of hard drives on servers and desktop PCs should pick up the rest. Software on desktop PCs should also be set to scan portable media in order to nab viruses that arrive, not through the network, but on diskettes and CDs.

3.4.2 Keep up-to-date

To stay effective, antivirus software needs to be constantly updated with the signatures of recently discovered viruses. This chore, fairly easy if you only have a single PC or a small network, can become an administrative nightmare in a large network, especially if there are remote or laptop users.

Many makers of antivirus software enable you to automatically update antivirus software over the Internet. Although the savings in labor costs can make this an attractive option, this option leaves decisions about what to install on your network entirely to the antivirus software vendor, so you may want to keep a human in the loop.

3.4.3 Patch holes

Modern software is very complex, making it difficult to thoroughly test for security holes. Often these holes are discovered after software has been out for a while. At this point, the vendor will release a software patch, usually available on its Web site. Many computer break-ins can be prevented simply by keeping your software patches up-to-date. Regularly schedule a check of software patches issued by your software vendors and use them where needed. Remember, do NOT install software patches that arrive unsolicited through your e-mail. These are worms.

3.4.4 *Use adware eliminators*

This is a specialized kind of software that helps you track down adware and spyware, and block those annoying popups. This is not the same thing as antivirus software, although some spyware is actually a virus and will be nabbed by your antivirus software. You can get adware and popup blockers for free on the Internet, but be careful what you download so you don't actually bring in more adware. As with antivirus software, no software will eradicate all offenders, making it a good idea to run more than one kind.

3.4.5 *Don't use file-sharing programs*

If the dubious legality of file sharing doesn't faze you, consider this: File-sharing programs such as Kazaa and Grokster are major culprits when it comes to dumping adware and spyware on your computer. They usually inform you in a disclosure box that they are supported by—and will install—adware. Read this box carefully before you click on the “agree” button—you may be agreeing to as many as ten different adware programs. Some even install spyware that they tell you nothing about. This adware can be very difficult or impossible to remove from your system. If you remove the adware, the file-sharing program may simply reinstall it or the file-sharing program may not work without the adware. Occasionally, the file-sharing program itself is a spyware program.

3.4.6 *Screen your cookies*

Many Web sites leave behind markers or cookies in your system. These cookies are used by Web sites to track you—for instance, to tell if you're a repeat visitor. Your browser preferences can be set to alert you with a dialog box before it accepts a cookie. This enables you to only accept cookies from Web sites you want to be known to.

3.4.7 *Some miscellaneous security precautions*

- Don't give out more information than you have to when you register software or buy something on-line.
- Turn off computers when you're not using them—PCs can't be hacked when they're off.
- Some e-mail programs have multiple security holes making them easily infected. Don't just use the e-mail reader that came with your operating system—shop around for the most secure program you can find that also meets your needs.
- Disinfect PCs with antivirus software before connecting them to your network.
- Scan all removable media.
- Disable Java and JavaScript. This will limit your interaction with some Web sites but protect you from malicious scripts.

3.5 *Broadband routers — simple, cheap firewalls, sort of*

An inexpensive but effective way to add a layer of rudimentary firewall protection to a small network with a broadband Internet connection is to use a broadband router instead of a switch to connect to your cable modem or DSL modem. The router provides an important buffer between your network and the Internet through the use of Network Address Translation (NAT). NAT enables several network users to share a single IP address and also provides an important security buffer for your network by hiding your real IP address from the Internet.

An Internet Protocol (IP) address is a 32-bit number that identifies an Internet host. IP addresses provide universal addressing across the Internet. IP addresses are placed in the IP packet header and are used to route packets to their destinations. Because IP addresses are difficult to remember, most also have text equivalents such as `whitehouse.gov` or `blackbox.com`. A database program called Domain Name Service (DNS) keeps track of the names and translates them into their numeric equivalents.

NAT assigns a local IP number to each PC or other network device and then maps local LAN IP addresses to your Internet IP address and vice versa. This masks your internal IP addresses from the world, making it far more difficult to hack into your PCs.

3.6 *Firewalls*

A firewall controls traffic between two networks. The most common application for a firewall is to control traffic between a private network and the Internet in order to intercept outsiders trying to break into the private network. A firewall exerts this control by applying rules to information—primarily IP addresses and port numbers—found in incoming network packets.

The word firewall is most often used to describe a freestanding firewall appliance that provides intelligent, port-based security, although some low-end firewalls are software based. Often services, such as NAT, that are provided by a broadband router are also described as being a firewall. Although the primary purpose of firewalls is to protect your network against unauthorized logins from outside your network, they're also sometimes used in companies or schools to prevent users from going out to the Internet. For instance, you can set the firewall so users can use e-mail but can't browse the Web.

3.6.1 *How firewalls work*

Computer services are based on ports. These ports are not the holes in the back of your computer that you plug cables into. These TCP ports for IP services are abstract, logical connections in your computer that enable it to handle multiple applications over a single network connection.

Every computer has 65,000 ports. Which ports provide which services is established by convention. For instance, Web browsing and HTTP traffic generally goes to Port 80, e-mail traffic uses Port 25. Your computer looks at which port data is entering to decide how to treat it.

Because there are only five commonly used Internet services, computers have thousands upon thousands of unused ports. These ports can be either open or blocked. Applications open ports to connect to services—for instance, when you start an e-mail program, it opens Port 25 to “listen” for SMTP traffic. Many applications open ports without your permission or knowledge.

Hackers probe computer networks by experimentally sending code to different ports. They’re looking for responses from open ports—the unused ports that your applications left open behind your back. When a hacker finds an open port, he’s found a way into your network.

Your goal is to make sure that unused ports are blocked and that your network only accepts legitimate requests for service. This is where a firewall comes in.

The firewall blocks unwanted traffic while letting through the traffic you want. It makes decisions that allow or deny access to services and ports on your firewall.

A firewall enforces your access control policy, but it’s up to you to decide what that access control policy is. You can block whole ranges of ports—everything that you do not require to be open. Firewalls generally come preconfigured to deny all access to all ports. It’s then up to you to instruct your firewall to allow network traffic through to specific ports on specific PCs in your network. The trick is to set it up to allow only necessary traffic through to the right ports, leaving all unused ports protected by the firewall.

When a request for a service is made, the firewall inspects the request to make sure the type of request matches an available port. Only traffic for advertised services is allowed through the firewall—all other traffic is dropped.

The firewall hides computers from the Internet altogether if they don’t provide services to Internet users. For instance, if your only connection to the Internet is for e-mail, the firewall can be set to only let e-mail traffic through to the appropriate server, shielding the rest of the network from the outside.

Many firewalls today offer protection against denial-of-service attacks, identifying and blocking any address trying to flood your Web site with traffic. However, hackers have learned to get around this by launching a DDoS attack in which they hijack other computers or spoof addresses, making it look like the attack is coming from many different IP addresses. Your options to defend against these attacks are extremely limited. One step you can take is to request that your ISP apply egress filtering to make sure that their outgoing packets carry only IP addresses that belong to that ISP. You may also scan for and remove DDoS clients, so that your PCs aren’t used as attack tools. These steps, unfortunately, help to protect other sites from attack but don’t do much for your immediate problem. However, if everyone scans their PCs for DDoS clients and more ISPs implement egress filtering, it can help to make life more difficult for hackers trying this method of attack.

3.6.2 *Keeping track of your firewall*

A firewall is valuable for its logging and auditing functions, providing summaries about what type and volume of traffic passed through it, and also what kinds of break-ins were attempted. When you check firewall logs, you'll find that network probes are surprisingly common—the Black Box network firewall records thousands daily. The logs will show you where hackers are trying to break in. You should examine logs on a regular basis—preferably once a week—and adjust the firewall accordingly. Some ports are favorites with hackers and you should pay special attention to attempts to access these ports. You should also regularly scan your network to find open ports. Block any unused ports you find open.

3.6.3 *What a firewall doesn't protect against*

As important as a firewall is to protect against hackers, it's important to remember that it does NOT provide complete network protection.

A firewall can't protect you from traffic that does not pass through the firewall, for instance, sabotage from inside your network or corporate secrets leaked out on floppy disks.

A firewall generally doesn't guard against viruses that typically piggyback on what looks like legitimate traffic. Many firewalls are sold as offering virus protection but only filter out some common types of viruses. A very few high-end firewalls offer complete virus protection but generally you must run separate virus-protection software in addition to your firewall.

3.6.4 *Types of firewalls*

A wide range of firewall options is available for applications ranging from a single PC with a broadband Internet connection to a large corporate network. Firewalls can be software that runs on a PC or a separate hardware device that has built-in firewall software.

Software firewalls have the advantage of running on a PC so you don't need to buy any extra hardware. Another advantage is that they're generally inexpensive—some good basic firewall protection can even be had for free. The downside of software firewall protection is that it's really only suitable for small applications such as protecting an individual PC or home network with a broadband connection. Also, because these firewalls are software that reside on your hard drive, they can be corrupted.

Hardware firewalls can refer to minor firewall capabilities, such as NAT, built into a broadband router, but more often they refer to a specialized firewall appliances that provide intelligent filtering and logging services as well as NAT and basic virus protection. These dedicated firewall appliances generally provide more protection than software firewalls and are more suitable for large networks. Plus, because their operating system is permanently encoded in hardware, they're more difficult to take down than software firewalls are.

Packet filtering or network layer (Layer 3) firewalls make decisions based on the source and destination addresses and ports in IP packets. This fastest and simplest form of firewall protection is really no more than a simple sorting algorithm. Generally these firewalls enable you to have some control through the use of access lists. Packet filtering can also often be performed by other network devices such as routers and is generally what you get when you download free firewall software.

Packet filtering works well for small networks but when applied to larger networks, it can quickly become very complex and difficult to configure. Packet filtering also cannot be used for content-based filtering and cannot, for instance, remove e-mail attachments. This type of firewall often has little or no logging capability, making it difficult to determine if it's been attacked.

The more sophisticated *proxy or application layer firewall* deals with network traffic by passing all packets through a separate "proxy" application that examines data at an application level.

A proxy firewall doesn't allow a direct connection between your network and the Internet. Instead it accepts requests and executes them on behalf of the user. For instance, if you're on network behind a proxy firewall and type in <http://www.blackbox.com>, the request goes to the firewall, which gets the page on your behalf and passes it to you. This process is transparent to users.

This proxy system enables you to set a firewall to accept or reject packets based on not only addresses and port information but also on application information. For instance, you can set the firewall to filter out all incoming packets belonging to .EXE files, which are often infected with viruses and worms. Proxy firewalls generally keep very detailed logs, including information on the data portions of packets.

Proxy firewalls are slower and require more hardware than packet filtering, however their greater versatility enables you to enforce tighter security policies.

When a firewall is described as having *stateful inspection*, it means that it examines packets at the network layer like packet filtering does but, rather than just applying simple filtering rules to this information, it uses it in an intelligent way to block out unauthorized traffic. A stateful inspection firewall analyzes data to make sure connection requests occur in the proper sequence. These firewalls track each communications session from start to end and enforce set rules based on protocol, port, and source and destination addresses. By maintaining all session data, the firewall can quickly verify that new incoming packets meet the criteria for authorized traffic. Packets that aren't part of an authorized session are rejected.

Stateful inspection firewalls have the advantage of being both smart and fast.

Packet-based, proxy, and stateful inspection used to be distinctly different types of firewalls, but today nearly all modern firewall appliances provide packet-based, proxy, and stateful inspection firewalling.

3.6.5 *DeMilitarized Zone (DMZ)*

Your network may include Web servers, FTP servers, or mail servers that are regularly accessed from the Internet. These servers, by their very nature, need to be open to the Internet. But you don't want the computers in your internal network open to the Internet and vulnerable to hacking.

For this reason, it's a good idea to isolate public servers from your private network in a DMZ. A DMZ is a separate network that's home for your public servers.

Both users from the Internet and users from the secure network may access servers in the DMZ. Traffic may not travel from the Internet or DMZ directly to the secure network without first going through a proxy server (usually a firewall appliance).

3.6.6 *Dual-homed gateway*

An easy way to build a DMZ into your network is through the use of a dual-homed gateway. A dual-homed gateway is usually a firewall, although it may also be a router. It features two LAN ports—one for your secure network and one for your "public" network—as well as a WAN port for connection for the Internet. One LAN port connects to the DMZ; one connects to your secure network. Each port has a different security policy, allowing a different type of access. The firewall automatically directs traffic to the appropriate port by IP address.

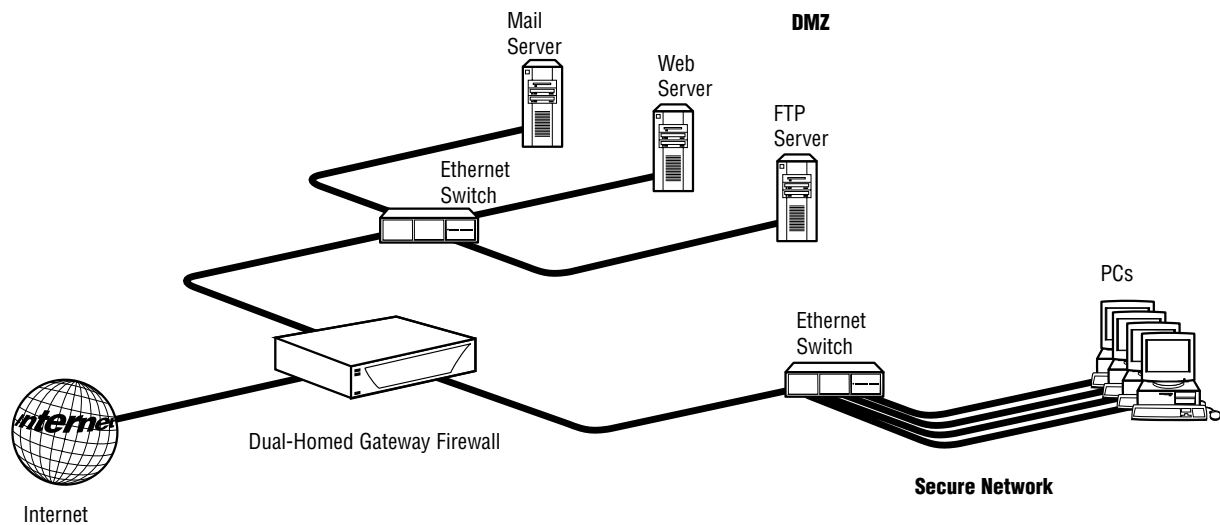
The dual-homed gateway serves as a secure point of control between these two network segments and the Internet, intercepting all traffic between the Internet and your secure site and screening all services and access through proxy servers on the firewall.

The dual-homed gateway firewall has the advantage of being both simple to set up and very secure. Many firewalls are specifically designed to support the dual-homed-gateway configuration and offer a convenient way to set up a DMZ, even in a very small network.

Dual-homed gateway firewalls tend to be rather inflexible, but they have the advantage of being easily managed by nearly anyone with basic networking skills. This makes them a popular choice with administrators of small-to-medium-sized networks who must have an e-mail and a Web server open to the Internet while at the same time protecting a secure, private network.

Other, more complex firewall configurations, such as screened subnet, offer more versatility and are more suitable to large networks but require a great deal of expertise to implement effectively.

Simple Dual-Homed Gateway Application



3.7 Backup and recovery

It's easier and less expensive to prevent problems before they happen, but even in well-defended networks, you have to assume that eventually the unthinkable will happen—a virus will get through and spread among your PCs or a hacker will break into your system and destroy files. And remember, there are other, more mundane dangers to your network—for instance, equipment failure, flood, or theft—that can cause at least as much damage as a virus infection. With a bit of planning, you can contain problems and restore network CPUs quickly.

Back up your files on a regular basis so if your network is invaded, you can replace corrupt or infected files with your backup copies. With regular nightly backups, the worst virus infection will never cause the loss of more than a day's data. Backup copies should always be stored on hard media in a separate location—NOT on a server connected to the network.

Have a plan to cover unexpected disaster. In the case of a virus infection, you should have a clear plan for disinfecting PCs and restoring data. If anything is absolutely mission critical, you should even have a plan in place for replacing hardware that goes offline for purely mechanical reasons.

Eradicating viruses and restoring lost data may, unfortunately, be the easiest part of recovery. What can be more difficult to recover from is lost credibility and breaches of confidentiality. For this reason, prevention is always the first goal; recovery is the second goal.

4.0 Conclusion

Every network administrator is faced with this dilemma: the Internet can be a dangerous thing for your network, but your network needs to be connected to it. What you can do is reduce your risk to an acceptable level by implementing the following steps:

- Have a plan.
- Use common sense and know what to avoid.
- Keep viruses and other nasty critters at bay with the appropriate software.
- Protect Internet-connected networks with a firewall.
- Back up data on a regular schedule.

5.0 Useful links

Viruses

<http://www.symantec.com/>

<http://www.mcafeeasap.com>

Spam

<http://www.spammotel.com/>

Hoaxes

<http://www.snopes.com/>

Adware and Popups

<http://www.cexx.org/adware.htm>

<http://www.lavasoftusa.com/>

<http://www.popup-killer.info/>

Firewalls

<http://www.blackbox.com/> (search firewall)

<http://www.zonelabs.com>

<http://www.robertgraham.com/pubs/firewall-seen.html#1.1>