Global Information Assurance Certification Paper

Are you a good Internet neighbor?

Scott Lawler
October 24, 2000

The Internet is a dangerous place...for technical as well as legal reasons. Internet connections have the potential to significantly revolutionize business, increase efficiency, and provide access to vast amounts of information. However, failing to manage the additional risk could be very expensive. Criminal and civil judgments are coming to those who fail to police themselves. Many networks remain poorly configured and permit malicious activity to continue which often affects others. Harboring malicious Internet activity, knowingly or unknowingly, could add significant technical, financial, and legal risk to your organization.

The neighborhood:

To begin, consider an analogy representing the Internet as a neighborhood. Each neighbor is expected to maintain some semblance of order and discipline on their property by a combination of legal and social influence depending on the local area. As an example, a malicious guest visits your home and runs through your entire neighborhood phone list calling each neighbor and letting the phone ring only once. This description is very similar to network mapping or Internet Control Message Protocol (ICMP) scanning--except that you have millions of next door neighbors on the Internet. Debate rages over whether or not ICMP scanning is malicious activity or not. ICMP is an important troubleshooting tool as well as an essential protocol for network operation in many cases. However, when ICMP is misused it can become an annoyance or a communication mechanism for trojan horses. Your neighbors quickly learn to ignore that one-ring phone call or those random ICMP packets from all over the globe. Stealth scanning is also gaining popularity. Many attackers will send crafted packets that will not make a proper connection. These stealthy techniques craft odd packets to pass through firewalls and help identify remote operating systems. Why would someone need to know what operating system version your computers are running? Exploits to break into just about every operating system are available to everyone on the web. Are crafted packets dangerous? Just what kind of lock is that on your front door? Deadbolt or not? Do you lock your interior doors? Are your guests conducting research prior to committing a crime?

The crime:

When does scanning become worthy of civil or criminal action? Ringing one telephone once is an annoyance. It doesn't cost the user of the telephone any more money or time. If a caller calls a million homes, is that still just an annoyance? In this case, the caller consumes a significant amount of bandwidth and telephone switch processor time. Think about Mother's Day and the millions of calls made that day. At some point, Internet scanning becomes theft of service. The malicious source is denying legitimate users the use of a resource. ICMP may only be a small fraction of the overall traffic volume on any single link. But, when a million links are considered, the theft of service multiplied across many service providers and many customers is significant. This caller is

using your telephone to commit a crime under US law. The Internet connected network in your backyard, that you own or manage, could soon be part of a criminal or civil action.

Port scanning is interesting from a legal perspective. In this case, the guest from your network backyard is running through the neighborhood trying all the doors and all the windows to see if they are open. Everybody's doors. Everybody's windows. Remember that you have millions of next door neighbors on the Internet. Your guest is conducting precursor to attack activity...looking for an easy place to break into. Of course, all your neighbors lock all their doors and windows. On the Internet we all know this isn't the case. Many of your neighboring networks are easily broken into. Your guest is about to become a criminal using your systems. The legal question is if your guest, or network user, is breaking into someone else's computer are you liable from a legal perspective? If that same guest leaves your home drunk, driving under the influence of alcohol, there are clear legal implications for you as the host if that guest injures someone.

Computer intrusions should be considered in this neighborhood analogy as well. Your guest throws firebombs over your back fence into the yards of your million neighbors. Most of your neighbors apply some diligence and will simply ignore the small flame, watch it burn on the sidewalk, or put it out promptly. Some of your neighbor's homes will burn. If you have a user on your network running attack or exploit scripts actively breaking into other computer systems, you are witnessing criminal activity. Are you now an accomplice if you permit this activity to continue? Should you be partially liable for failing to control the guest in your backyard? A similar analogy may hold for users releasing malicious code on the Internet from your network. Did you have knowledge of this activity? If you did, what did you do? Conversely, do you protect your inbound links from malicious code such as email viruses? Why do you allow malicious code to propagate out of your network? From a legal perspective, the telephone providers are not responsible for the content of the telephone calls. However, misuse remains a crime and should be reported. One form of abuse is called "wire fraud" under US law.

Denial of service attacks are a significant threat to the availability of Internet services. Continuing the analogy, your guest is now throwing nails in your neighbor's driveways and the streets surrounding your home. Some of your neighbors have wide driveways and simply ignore the nails. Other neighbors will see traffic jams or have a flat tire miles (many router hops) away and not know why. The victim of the attack will see packets coming from all over with spoofed addresses. Basically, the victim's home is under assault from all directions...so much so that they can't even see out. Even very large bandwidth connections can be completely shutdown with a few hundred well connected Denial of Service agents. These agents (sometimes called daemons) are your internal computers that are now running malicious attack software controlled by outsiders. These compromised computers in your network can now attack someone else. The current Distributed Denial of Service (DDOS) tools contain a variety of attack methods, stealthy communication, can spoof source IP addresses, and are quite flexible and easy to use. Basically if you allow an outsider to compromise a computer in your network and attack someone else, you could be an accomplice to a crime. For failing to apply due diligence, you are now potentially liable for civil and possibly criminal actions. The damages can easily add up to millions of dollars a day for a busy Internet

commerce site.

Criminal activity is not always obvious. What are are your liabilities if a guest of yours stores and distributes exploits and malicious code on your systems? Our favorite guest basically stores bomb making materials in your garage. The guest may also let a million of your closest neighbors know where these tools are through chatting with them or website links. Some of these exploit tools are also network management and system analysis tools. The right tools can be quite effective for securing systems when used by system administrators. In the hands of a malicious guest, the same tools can be used to exploit vulnerabilities in computer systems around the world. Are you liable for these tools being stored on your systems? What is the legal difference between a system administration tool and a hacking tool? Many hacker web sites proclaim they distribute information on how to hack solely for system admins to learn to tighten up systems. Yet, hackers use the same sites to learn to break into your neighbor's systems. This dichotomy remains a tough one for network security professionals. They are often faced with the dilemma of publishing a new vulnerability with no fixes or waiting and hoping the vendor will respond quickly. Education and software maturity across the Internet over time may be viable answers.

Protect yourself:

What should you do? From a technical perspective, there are a number of issues to consider but breaking them down into people, process, and technology will help make the issues more manageable. First, the system and network administrators must be knowledgeable and security aware. Operating system, application, and security training are essential. Sufficient prioritization and time devoted to network and host security will significantly reduce the risk of your network being compromised and used to attack others. What are your critical systems? What is the value of the information on those systems? What does normal operation look like? Just keeping an eye on your backyard and understanding what some of these malicious activities look like will help you react quickly and early enough to significantly mitigate risk.

From a process perspective, a tight security policy is critical. Your guests should know and understand what is acceptable behavior and what is not. Many organizations have sound policies but fail to enforce them. Or, enforcement consists of simply disabling an anonymous dialup account. Attackers will just use another free disk to get right back on. There is a tangible cost associated with monitoring for malicious activity. Many network managers do not devote resources to this task because they are not aware of the financial risk associated with permitting malicious activity. These financial risks will become more apparent as the frequency of civil cases increases and judgments rise. The insurance industry is already considering network security as a factor in risk mitigation. Your organization must have an established abuse policy, process, and corrective action plans. Auditing is important not only for external users but internal users as well. You need to know when you have a suspicious guest, employee, student, visitor, or unauthorized user. Audits must be reviewed periodically to stand up in court. These audits must be documented. Initials and signatures in paper log books are excellent. Another important process is handling abuse email. Some organizations do not respond to abuse email. That corporate decision is acceptable provided

the complaints are routinely acted upon. However, responding to abuse complaints should be a common professional courtesy. Many organizations will respond with brief sanitized messages removing any internal proprietary details. Professional CERTs will not publicize any information that you send them. However, if you are involved in a major intrusion or denial of service, it's important to have an organizational process in place to handle information requests from the media...soon to be followed by warrants for technical documentation such as log files and system images. Legal and CIO or CEO reviews may be needed depending on the severity of the situation. Established processes are important.

Defense in Depth:

Technically, there are well known actions that you should take to manage, control, and monitor your network. Key concepts are available in most Defense in Depth documents and textbooks. We'll focus on bare minimums to mitigate the most likely technical and fiscal risks.

Defense in depth minimums:

- default deny all-permit only what you need policy
- intrusion detection
- host protection
- protection zones
- ingress and egress filtering

First, you should deny all traffic inbound that you do not need for a known legitimate use. A security policy is essential for identifying what these services are. Intrusion detection and host protection are also important as another layer of alerting. You need to know when malicious activity is happening--what you don't know CAN hurt you financially. Under US law, a network service provider can monitor activity for "the protection of the rights or property of the provider or the service". The concept of using DMZs is well known in the Internet industry however many organizations don't take advantage of them. A DMZ puts your high risk systems outside your protected internal network. Tightly configure and monitor DMZ hosts for malicious activity. Daily log file reviews are essential. Similarly, user groups can be partitioned off in logical or physical groups that use similar services. These groupings could prevent a malicious user in one zone from compromising systems in another zone without being detected.

Ingress filtering is fairly obvious. You know you want to keep malicious Internet users out of your systems and most organizations block unused inbound traffic. However, do you periodically double-check those configurations? Often a test or temporary rule gets left in. Or, a legitimate hole is left open for a system that is no longer in place. Finally, egress filtering is a very important concept to reduce legal liability. Egress filtering means that you only allow out traffic that originated from inside your network. Only your email server can send mail out. Only your web server can service http requests. Only clients you know about can access internal and external network services. Intrusion detection on the inside of your network can also act like an internal

burglar alarm to let you know when your systems are at high risk. Egress filtering limits what your network hosts can do...it also limits what a malicious guest can do with your systems. Many DDOS tools rely on spoofing IP addresses to hide where they are coming from. These are VERY noisy devices and simple auditing on outbound traffic could let you know when you have internal security problems. There are many excellent technical resources on Defense in Depth and egress filtering available. Overall, technically configuring your network to first prevent intrusion in the first place and secondly to limit that damage an intruder can do will demonstrate due diligence and significantly reduce your risk of technical compromise and the potential of facing significant legal liability.

Large networks:

From an enterprise perspective these abuse challenges are much more difficult but not impossible to manage. Larger National Service Providers (NSPs) for example cannot effectively route only internal traffic. In the neighborhood analogy, the larger NSPs are like the highway system connecting neighborhoods. In this case the NSPs must filter what is practical at the enclave perimeters and monitor core hubs and gateways for anomalous activity. NSPs are less liable for malicious activity because they are farther disconnected from the activities of individual users. Regional and small Internet Service Providers along with most enclave network managers bear the burden of malicious user monitoring. The only way to reduce risk across the Internet is for local network managers to tell malicious users "Not in my backyard!" However, NSPs do have a significant role in helping to identify where the criminals are really coming from. Locating attackers takes a team effort between the victim, the victim's ISP and often several other Naps and/or ISPs along the path. Collaboration is essential to reduce the risk of relying solely on Internet connectivity for business or government operations. As a best practice, NSPs and ISPs should have rapid access to 24/7 POCs for all upstream and downstream network connections. Contacts should strive to quickly exchange log file excerpts. Rapid information exchange significantly aids identification and tracking of attackers. As long as significant numbers of networks remain unprotected, the risk of extended Denial of Service remains for all Internet Service Providers and users.

Due diligence:

Just exactly when does the potential legal or punitive damage risk start? Due diligence is a not a matter of buying the latest security product every month. Due diligence is applying well known security fixes to systems and continuing to improve security posture over time. From a very basic perspective here are the must do items to demonstrate basic due diligence:

<u>People:</u>

Provide sufficient staff to handle the systems in operation

Train the staff - in critical system operating systems, applications, and security

Process:

Implement a signed Security Policy

Document periodic audits and reviews

Maintain user accounts - watch for new unauthorized accounts, promptly disable, and control access

Technology:

Implement a tight perimeter policy - router access control list or tight firewall configuration (Simply having a firewall is NOT enough.)

Update operating systems and application versions and patches

Install host protection and monitoring on critical hosts - tcp wrappers and host monitoring software

Install and monitor Intrusion Detection Systems - regular reviews of the logs are critical

Professional networks:

Keep in mind the above list is a bare minimum to help keep your network from attacking others. If you do these things your network will probably be passed over by the lesser skilled attackers. However, a CIO supporting a professional system administration team will continually review procedures and implement new best practices from the Internet community. A professional team also monitors many different devices in Fault, Configuration, Accounting, Performance, and Security areas. For example, professional network managers conduct internal and external vulnerability scans to look for and patch holes before an attacker exploits them. Professional teams tightly control system configurations. Professionals also baseline, monitor, and configure system alerts to know when unusual activity occurs. Skilled network managers will use baselines to establish rate limits on key network devices. Teams with mature processes make backups but also store a backups securely off-site and restore from the backups periodically to verify the process. The challenge to these professionals is to share your best practices with the Internet community to help others improve security as well. We need to help our neighbors to reduce the risk for all.

Negligence:

What is obvious negligence? If a network configuration allows virtually all connections in and out, has unpatched hosts, no effective monitoring, no response to abuse, and has repeated compromises, the network manager/owner would be considered clearly negligent from a people, process, and technology perspective. If your systems are used to repeatedly attack others, you could be liable in a criminal or civil case. These judgments will be higher and higher as more and more companies rely on Internet commerce as a significant revenue stream. Negligence is leaving that

malicious guest alone in your backyard...to wreak havoc on your neighbors unfettered.

Conclusion:

To wrap up, clean up your network and continue updating and monitoring to keep it clean. If you perform basic network security functions they will significantly reduce the technical and legal risks associated with connecting to the Internet. Using the backyard analogy may help communicate these concepts to non-technical people. For further information, see the references.

Areas for further research:

At what point does scanning become a denial of service?

What are the implications of Internet abuse crossing the globe through a wide variety of legal systems?

From a legal perspective, what is the definition of "due diligence" in rapidly changing Internet security?

Is the owner of a compromised home computer legally liable for attacks originating from that system?

How long should an ISP maintain audit records? If no records are available, there is no evidence...

==========

References:

18 U.S.C. Sec. 2510-2522. Electronic Communications Privacy Act of 1986.

47 U.S.C. Sec. 223. Obscene or harassing use of telecommunications facilities under the Communications Act of 1934.

Stevens, W. TCP/IP Illustrated, Volume 1. Addison-Wesley. (1994.)

CISCO White Paper. "Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks." (February 17, 2000.) URL: http://www.cisco.com/warp/public/707/newsflash.html

Carnegie Mellon. Software Engineering Institute. CERT Coordination Center. "Protect your Web server against common attacks." (12 June 2000) URL: http://www.cert.org/security-improvement/practices/p082.html

Dittrich, D., Weaver, G., Dietrich, S., Long, N. "The "mstream" distributed denial of service attack

tool". (May 1, 2000.) URL:
http://staff.washington.edu/dittrich/misc/mstream.analysis.txt

Zwicky, E., Cooper, S., Chapman, B. Building Internet Firewalls. 2nd Edition. O'Reilly. (June 2000.)