Testimony of Gregory Neal Akers
Senior Vice-President and Chief Technology Officer
Government Solutions Group and Corporate Security Programs
Cisco Systems, Inc.

Hearing Before the

House Committee on Government Reform

Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census

September 10, 2003

Chairman Putnam, Ranking Member Clay, and other Distinguished Members: Thank you for the opportunity to testify today regarding protecting the nation's computers against the growing threats caused by worms and viruses. We are enormously dependent on the correct operation of the Internet, and recent surveys show that Americans are concerned for the safety of business conducted via the Internet.[1]

My Background

My name is Greg Akers, and I am Senior Vice-President and Chief Technology Officer for Government Solutions and Corporate Security Programs at Cisco Systems, Inc. In addition to my present executive responsibilities, I have held senior technical positions at Cisco, including network engineer and vice president of our Technical Assistance Center (the Cisco TAC). Additionally, I am a Cisco Certified Internetworking Engineer (CCIE #1037). Prior to joining Cisco, I spent fifteen years designing, building, and running large networks for "Fortune 100" companies. In 2002, I served as the President of the IT-Information Sharing and Analysis Center (IT-ISAC) and as the Vice President in 2001. Currently, I am a member of the National White-Collar Crime Board and the Board of Directors of the East Carolina Infragard.

Cisco and the Internet

Cisco Systems is the worldwide leader in networking for the Internet. Our networking solutions connect people, computing devices, and networks, and allow people to access or transfer information without regard to differences in time, place, or type of computer system.

[1] "The Internet and Emergency Preparedness: A joint survey with Federal Computer Week magazine", The Pew Internet Project, August 31, 2003, http://www.pewinternet.org/reports/toc.asp?Report=100

We provide end-to-end networking solutions that customers use to build a unified information infrastructure of their own, or to connect to someone else's network. An end-to-end networking solution is one that provides a common architecture that delivers consistent network services to all users. The broader the range of network services, the more capabilities a network can provide to its connected users.

Our core technology began with routers. Routers are what make the Internet work. They act as multi-protocol translators that tie the disparate computer networks of the world together on the Internet, in much the same way that telephone networks in different countries connect and place calls to each other.

Cisco's success is inextricably tied to the Internet. Approximately 80% of Cisco customer support calls are resolved over the Internet. In addition, we estimate that about 85% or more of sales of Cisco's products and services are completed via our website, cisco.com. Therefore, we are very concerned by worms and viruses that threaten the correct operation of the Internet. The Internet is "mission-critical" to Cisco's business.

In my brief time with you today, I will address worms, viruses, and vulnerabilities, as all three are tightly integrated. I will describe issues around vulnerabilities, how vulnerabilities are discovered, and Cisco's process for managing product security incidents, including how we disclose vulnerability and remedies to customers. I will also describe some techniques to reduce the threat of these vulnerabilities.

Vulnerabilities as Vehicles for Viruses and Worms

Viruses and worms exploit a vulnerability to propagate; therefore we will treat viruses and worms identically in this discussion. For the purpose of this testimony, we will focus on vulnerabilities, which we define as a set of conditions that leads to implicit or explicit violations of the confidentiality, integrity, or availability of an information system. Examples may include any one of the following actions performed without authorization:

- Executing commands as another user;

- Accessing data in excess of specified or expected permission;
- Posing as another user or service within a system, or:
- Causing a denial of service.

As more business is conducted using interconnected information technology, the risks of these systems to various attacks is also increasing. The type and scope of such threats can change daily. Additionally, threats are becoming more covert and intricate, which makes them harder to track, root out, and identify.

How are Vulnerabilities Discovered?

Vulnerabilities are uncovered in a variety of ways, such as by vendors during testing, in the course of normal customer use, by vendor-neutral security organizations conducting research, and by miscreants probing systems and programs.

Vendor Testing: As a vendor, Cisco regularly conducts extensive testing of its software and hardware to maintain and improve the security and stability of our products. As the latest vulnerability analysis tools become available or are developed internally, Cisco seeks to proactively identify enhancements and resolve issues, including a strong focus on security vulnerabilities. We consider a variety of factors, which can include the ease of exploitability, the critical nature of the service or protocol to the operation of networks, and the ubiquity of the equipment or application.

Customer Use: Many security vulnerabilities are discovered through customer use and are reported by way of a customer support organization. Vulnerabilities are not obvious as the root cause of a customer support case and may be difficult to identify as a true vulnerability. Customer in this context refers to any user.

Vendor-neutral Organizations: Vendor-neutral organizations, such as the Computer Emergency Response Team/Coordination Center (CERT/CC) at Carnegie Mellon University, coordinate responses to security compromises, identify trends in intruder activity, work with other security experts to identify solutions to security problems, and disseminate security improvement information to the broad community. Additionally, they maintain a database to provide early warning of vulnerabilities to Department of Defense (DoD) and other government users. In some instances, affected vendors may employ the assistance of a trusted intermediary such as the CERT/CC to coordinate a multi-vendor product security incident. This can be a valuable service, but it is dependent on the impartiality of the coordination center – If the organization becomes heavily reliant upon a government or commercial organization for funding, the trust placed in it by the community might be diminished to the extent that it can not operate effectively.

Miscreants: The miscreants who uncover vulnerabilities typically range from "script kiddies" (the cyberspace equivalent of vandals and hooligans), to professional "black hats" who work for organized crime, terrorists, nation-states, or some combination. While a "first-time" exploitation of a vulnerability may require some technical expertise, almost **anyone** can make use of exploitation tools afterward. Miscreants often publish these tools widely on the Internet and elsewhere. Many successful exploits are "only" a mouse-click away; no prior experience is necessary.

Public Notification of Vulnerabilities

A key to protecting our nation's computers is effectively sharing information about cyber threats, vulnerabilities, countermeasures, and best practices. Differing opinions exist regarding the most appropriate way to disclose vulnerabilities. Nevertheless, there appears to be little dispute that vulnerabilities should be disclosed in order to reduce the risks to information systems and to minimize or halt related malicious activity.

Vulnerability disclosure is not a simple process. Affected vendors must carefully consider multiple factors in light of the nature of the vulnerability at hand. When, for example, is the appropriate time to disclose? How much information about the specific vulnerability should be revealed? Should the disclosure be made to the public all at once time, or should certain entities, such as core internet service providers, receive some advanced notification before the vulnerability is fully disclosed to the public?

If vendors disclose vulnerabilities to customers and the public before fixed software or workarounds are developed and available, customers may face the risk that a miscreant will attempt to exploit the vulnerability. If the vulnerability affects systems in widespread use within critical infrastructures, the risk to national and economic security is magnified.

It is against this daunting background that a vendor, who seeks the best way to disclose a vulnerability to the public, must carefully determine how to best minimize the risks associated with the possible exploitation of that vulnerability during and after the disclosure process.

Cisco’s Vulnerability Disclosure Process: Cisco has long recognized the importance of disclosure of vulnerabilities, with a history of vulnerability disclosure dating back over a decade. In 1997, Cisco formally established its Product Security Incident Response Team (“PSIRT”), an internal, dedicated team of technical experts that handle the full scope of activities associated with handling vulnerabilities. The team members are selected carefully and are part of Customer Advocacy, Cisco’s customer support organization.

When the PSIRT team receives a report of a vulnerability, it researches the exploitability and scope of the vulnerability, and then attempts to fully characterize it. The team treats reported vulnerability cases very confidentially in order to minimize the risk of accidental leaks. Once the PSIRT team has made an initial assessment that a true vulnerability exists, it contacts the Cisco development teams who are responsible for providing a fix. While the fix is in development, the team will then determine whether and what kinds of pragmatic workarounds might be devised and deployed.

Once the fix and the workarounds are developed and tested, the PSIRT team carefully documents the vulnerability. Many factors are taken into account for the published web description of the vulnerability. Enough information needs to be provided for affected customers to protect their systems; nevertheless, certain key details are often withheld to prevent miscreants from rapidly developing malicious exploits.

The PSIRT team is responsible for the time when the associated security advisory and fixed software are posted on Cisco.com. The team provides information to other Cisco organizations who respond to inquiries from customers and others about the disclosed vulnerability. After the publication of the advisory, the PSIRT team solicits feedback from affected customers and researchers to help monitor the effectiveness and viability of the fix provided. Based upon such ongoing post-disclosure monitoring, the team will continue to periodically post updates to the security on Cisco.com until the threat of an exploitation of the vulnerability has been successfully thwarted.

Mechanisms that Exist for Protecting Systems

Web traffic and mail are the two most common transport mechanisms for viruses and worms. Code Red, Slammer, Blaster, Nachi, SoBig, and most other worms and viruses entered networks through services that were specifically permitted. A typical network is expected to permit e-mail, web browsing, and news service between internal and external systems. Understanding this opportunity, attackers seek obscure ways to send their own data into the network mixed in with the normal traffic destined for web browsers, e-mail clients, and news readers.

There are many defense mechanisms designed to help protect networks and host systems from the threat of viruses, worms, and direct attack. However, such mechanisms are limited, both by their design and by the skill set of the person who configures them.

Properly configured and maintained firewalls can protect a network from an attacker trying to directly access the network from the outside. However, a firewall used alone lacks defense in depth, and cannot reliably protect against all viruses and worms. In a common scenario, a firewall administrator may inadvertently open up access to a much larger range of network traffic than suspected while trying to solve an independent network communication problem through the firewall. When such attacks are active, it may only take moments for malicious traffic to travel past the firewall and infect vulnerable systems on the other side.

Virus Protection Programs: Virus protection programs exist for mail servers, the powerful computers which receive our mail from the Internet and sort them out for delivery to the end users. These programs regularly allow infected mail through because they have to sort through too many large messages and they can't handle the load. Even the most powerful servers depend on the e-mail administrators to keep their virus definition files up to date. For some large enterprise networks, it can take hours for the administrators to update the mail servers to catch the latest e-mail-borne virus, and that can only occur after the anti-virus vendor makes the latest definition files available.

Network Intrusion Detection Systems: Many network and system administrators rely too heavily – sometimes solely – on network Intrusion Detection Systems (network IDSes). These are devices that scan the traffic on the network and compare it against “signatures”, distinctive patterns of common attacks. IDSes are very good at detecting unusual traffic, but they should be part of a larger system for securing networked resources and not relied upon as a sole means of protection. Many newer viruses and worms are better able to disguise themselves as perfectly legitimate traffic, increasing the difficulty of identifying them as malicious traffic. An IDS is a warning device, providing indication that further action needs to be taken. IDSes do not block attack traffic alone. Appropriate actions must follow to respond to the threat.

Other Network Security Tools: Other tools exist that are not yet commonly deployed that may provide some added network security protection. These include tools that monitor the “flow” of traffic that travels across the network, and which then pass such flow data to a device like those made by Arbor Networks or Riverhead Networks for further analysis. These devices look at the larger view of network traffic and report anomalous behavior such as greatly increased traffic to a specific Internet port number, a typical pattern for a new worm. In a similar vein, Cisco offers a program called CSA, Cisco Security Agent to detect inappropriate attempts to access files and other unexpected system actions on a single computer or server. Unlike antivirus programs which wait for a specific, known virus to start attacking, these programs can alert system administrators before a new worm or virus can be identified, “fingerprinted”, and announced. These host-based solutions are not yet widely deployed, but do appear promising.

Today, there is no one right solution. Vendors, end users, and system administrators can benefit from further education regarding the value of multiple tools to effectively combat these threats. Presently, the only available solutions are reactive and time-consuming. Each class of tool presented above prevents some form of attack, and new tools are constantly in development.

Keeping Systems Up to Date

The deployment and ongoing maintenance of software patches, upgrades, and workarounds incur significant time and manpower costs. A network administrator may be faced with upgrading software or implementing workarounds on thousands of devices. In many cases, the administrator can not afford to simply reboot the entire network, particularly if the resulting interruption will interfere with mission-critical services. In addition, some service providers and similar organizations may have service-level agreements (SLAs) in place with their own customers who require pre-notification of maintenance. Some “customer’s customers” require maintenance to be confined to certain times of the day or strictly limit the number of maintenance events in a time period. Testing of software upgrades can be very time consuming. The demands on testing requirements have increased dramatically in the brief history of the Internet, some of it mandated by industry requirements, telecommunications regulations, and SLAs. Most network operators must contend with a myriad of testing requirements. Some testing is self-imposed because many networks are unique, and in today’s competitive network services marketplace, no one can afford to deploy new software without fully testing it in their own unique environment.

Another major issue is the potential complexity arising from even the simplest of vulnerabilities. Some vulnerabilities are resolved with a “one-line” change to the source code. Others might force a near-complete redesign of the system. Such severe changes can have a dramatic impact on the confidence level of the customer, particularly in mission-critical situations. Therefore, system and network administrators are very conservative about changing a working system, particularly to defend against a vulnerability that may have not been developed into a malicious exploit.

Vendors can help. The more painless they make the upgrade, the more likely users will implement deploy it. The less impact a patch has on a working system, the more likely the customer is to trust vendor. For example, most fixed releases for Cisco products are part of the normal development cycle, and contain additional fixes for a wide variety of problems plus the addition of new features. In some cases, where it is pragmatic to do so, Cisco releases software containing **only** the exact fixes necessary to close the hole. In some cases, customers are more confident with running such software and may validate it rapidly using a reduced testing regimen. The result is that fixed code can be deployed much earlier, minimizing the customer’s exposure to risk.

Vendors make every effort to release stable code, but often vulnerabilities are being fixed under severe time constraints. A miscreant might know about the problem and may be developing an exploit. At the same time, the product vendor is racing against the underground, trying to release a patch before the new exploit – possibly a new worm or virus – is released. Sometimes there's simply not enough time to test every possible combination of the new code if the vendor seeks to release the fix before the miscreants start attacking. Other times, a well intentioned researcher may indicate willingness to publish a vulnerability in a month. From the vendors view, a month might be enough time to write the fixed code, but not enough time to exhaustively test the fixed software.

The timing of vulnerability disclosure requires a fine balance of speed and quality. A blanket set of rules that define a timeline or a requirement may inappropriately force a vendor to release a fix before the software has been fully tested. If a software patch turns out to be unstable, end users and system administrators may decide not to upgrade. Yet, by not upgrading, the networks then may become susceptible to an exploitation of the vulnerability.

Ultimately, all of these solutions depend on humans to react and respond in a timely matter. Anti-virus software is useless against newer worms and viruses if the signature database has never been updated. Anomaly-detection systems such as network-based and host-based IDSes cannot react by themselves – someone has to respond to the alarms and mitigate the purported threats. Systems are not patched if security advisories go unread, or the fixed software is not downloaded and deployed, or customers can't figure out where to find security advisories and related fixed software, or researchers and customers can't determine how, and to whom, to report a vulnerability.

Summary

Our global infrastructures are interlinked in complex, sometimes little-understood ways, and some of the dependencies are surprising.

The global nature of the Internet means that no single country or industry group can address vulnerabilities in isolation. Success in this arena requires public-private cooperation. Our common goal is to reduce vulnerabilities, mitigate risks, identify strategic objectives, and share sound information security practices.

An example of a cooperative industry effort is underway within the National Infrastructure Advisory Council (NIAC). NIAC has a current effort to develop vulnerability disclosure guidelines that should prove useful for discoverers, vendors, users, and governments. The NIAC will also make specific policy recommendations for the President. The study will be available after it has been delivered to the President in the coming months.

National and economic security are forever intertwined. The industry leaders I work with understand their role and are willing to do their part to protect our national and economic security. I would like to thank you, Mr. Chairman and other subcommittee members, for inviting me here today. I am happy to answer your questions.