SECURITY

The Sad And Increasingly Deplorable State Of Internet Security, Revisited

Lisa Phifer and David Piscitello

Is your organization part of the problem, or part of the solution?

In the February 2003 issue of *BCR*, we claimed that, "overall, Internet security really is in horrible shape." We were convinced by computer crime statistics, incident reports and our collective experience that the security technology deployed to date had not proven effective. In fact, incident frequency and cost were increasing at an alarming rate, despite the fact that most organizations were claiming to have deployed state-of-the-art security defenses.

In our 2003 article, we also predicted that security would worsen before it improved. We cited insecure operating system (OS) and network architectures, lame authentication, poor software engineering, lax security management and creeping featurism as principal root causes for this "fall from security grace." We concluded with a measure of hope, however, suggesting that feature moratoria, software reliability agreements, administration improvements and perhaps more regulatory influence could improve Internet security.

Now, four years later, *BCR* has again invited us to comment on the state of Internet security. Overall, we find that while security threats have evolved, the root causes of security vulnerability haven't changed, and they are still being ignored in favor of "quick fixes" to ease security symptoms. Although most of our 2003 advice hasn't been taken, progress is being made in a number of areas, including more secure operating systems and protocols, unified threat mitigation and identity-based network access controls.

Change For The Better Or Worse?

In 2003, security incidents tracked by the Computer Emergency Response Team Coordination Center (CERT/CC) were already spiraling out of control. Automated attacks became so commonplace that CERT stopped counting them in 2004; instead, today's security metrics focus more on newly-discovered vulnerabilities (Figure 1) and their potential impact when exploited by attackers (Figure 2).

For example, in a single week in April 2007, high-impact vulnerabilities were identified in ARM mobile processors, Mac OS X hosts, and Oracle database servers, Microsoft DNS servers and Nortel VPN gateways. These vulnerabilities expose clients and servers to arbitrary code execution, privilege escalation, information disclosure and denial of service attacks.

Last year, three-quarters of the identified vulnerabilities were remotely-exploitable, meaning that they let cyber-criminals wreak havoc and steal data from a distance. Exploits for more than half of these vulnerabilities were available within just six days, but vendors averaged 47 days to issue corrective patches.

However, some aspects of threat mitigation have improved since 2003. Most businesses have adopted client antivirus (AV), and two out of three companies now use personal firewalls, often in conjunction with VPN clients. The number of emails carrying virus payloads has dropped from one in every 10 in 2004, to one in every 68 in 2006. Actual virus infections occur less frequently, and the impacts are now largely limited to unsecured home PCs. But even here we see progress, as an increasing number of ISPs now provide free AV for subscribers.

Unfortunately, viruses have been superseded by an onslaught of blended threats, including tenacious trojans and worms that implant stealthy key loggers and rootkits. Spam is worse than ever: more than 86 percent of the mail sent in 2006 was spam. Spam also facilitates malware distribution via botnets, collections of phished and pharmed hosts that have been duped into installing malware and/or disclosing financial and identity data. Peacomm is a peer-to-peer Trojan that typifies the insidious nature of today's malware: instead of being commanded by one easily-stopped master, this multi-headed beast lets an attacker run the entire botnet through any surviving peer (for more on botnets, see *BCR*, March 2007, p. 55).

Dave and Lisa own Core Competence, a network security technology consulting firm focused on emerging technologies and best practices. Dave also serves as a fellow on the ICANN Security and Stability Advisory Committee. They can be reached at dave@corecom.com and lisa@corecom.com, respectively.

Attackers in 2003 were like street thugs who used shotguns to blast as many victims as possible, seeking notoriety and retribution. By comparison, today's attackers are more like assassins: They are financially- and politically-motivated, more persistent and elusive, and more selective when choosing their targets.

In early 2006, distributed denial-of-service (DoS) attacks used large botnets to send spoofed requests to root and top-level DNS servers, attempting to overwhelm these select-but-critical Internet infrastructure components. Most of the phishing attacks in 2006 (84 percent) used financial brands (e.g., PayPal, eBay, Chase) to steal account holder identities and credentials.

Although text mass-mailing worms like Mytob and Netsky remain common, new worms like Stratio use randomized images to evade conventional spam detection methods. Image spam now comprises 31 percent of all spam and is largely sent for profit (e.g., black-market pharmaceuticals, stock price manipulation).

As systems, applications and networks evolve, so do vulnerabilities. In 2003, the major threats were viruses, insider attacks and laptop theft, according to the Computer Security Institute and the Federal Bureau of Investigation's annual survey. While these threats persist, companies now also worry about wireless abuse, Web application exploits and lost thumb drives. Companies also worry more today about digital copyright infringement, software licensing abuse and the difficulties of tracking sensitive (especially personal identifying) information.

Compliance pressures also are greater today than in 2003 for HIPAA, SOX and other regulations that force organizations to invest in measures to protect privacy and prevent private data leakage. However, when security budgets are focused largely on regulatory compliance, emerging technologies are more likely to be connected to corporate networks without adequate security assessment, hardening, policies or controls.

FIGURE 1 Annual Vulnerabilities Reported To CERT/CC

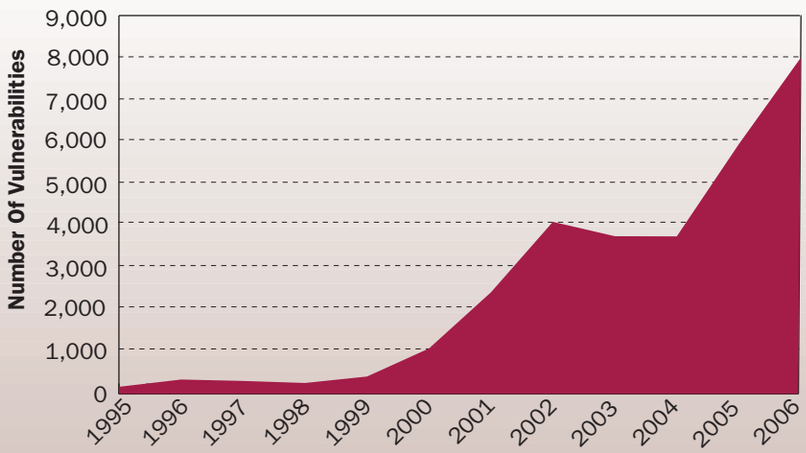

Source: www.cert.org/stats/cert_stats.html

We Are Still Ignoring Root Causes

Sadly, the vendors are still offering symptomatic pain relief. This has proven profitable, as companies and individuals are desperate for even a temporary reduction in visible, time-consuming activities like spam blocking. In the absence of more holistic alternatives, we continue to ignore many more serious threats, including:

■ **Vulnerable Routing Infrastructure**—The most widely used Internet routing protocols employ simple message digests to weakly authenticate routing peers. No cryptographic methods are currently used by routers to ensure that the routes advertised by peers are legitimate, or that those peers are authorized to forward those advertisements.

■ **Vulnerable Name Service**—The Internet's Domain Name System (DNS), which resolves website and email addresses into the IP addresses routers use, remains vulnerable to many forms of impersonation attacks. By impersonating a DNS resolver, an attacker can dynamically update or "poison" another name server with false records. By impersonating a name server, an attacker can falsify an entire domain's records by duping others into caching false entries, or accepting an entire false database during a zone transfer. These vulnerabilities can be used to abet phishing attacks and identity theft.

Solutions exist, in the form of RFC 2845 (Secret Key Transaction for DNS) and RFCs 4033, 4034 and 4035 which define DNS Security (DNSSEC), but these are not widely implemented yet (for more on this topic, see *BCR*, November 2006, pp. 44–47).

FIGURE 2 Criticality Of Security Advisories As Ranked By Secunia

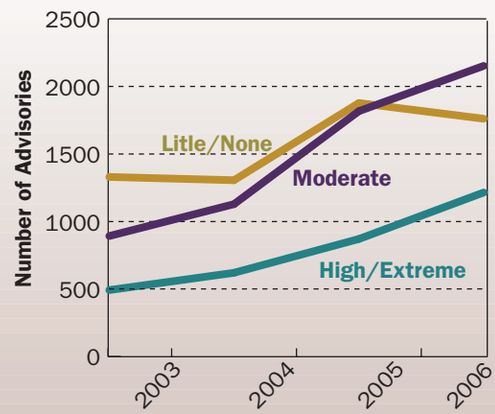

Source: secunia.com/advisory_statistics

We still innovate with the attitude that we can "add security later," even though that is always more expensive

■ **Weak Authentication**—We rely on text passwords, the least secure form of authentication. Nearly 80 percent of companies use password authentication, according to a recent Infosecurity Europe survey. The typical enterprise employee has 4 to 20 passwords, and many use the same passwords for work and personal logins. While we see a modest increase in the adoption of multifactor and biometric authentication systems, many organizations are content with plain-text passwords, spending their budgets instead on password management systems to reduce help desk costs.

■ **Lack of Source Verification**—We have made little progress in curbing impersonation. IP spoofing (forging the source address of IP packets) is arguably the most prevalent form of impersonation, used by many Internet attacks to make isolating and squelching the attack more difficult. For example, in DNS distributed denial of service (DDoS) amplification attacks, also discussed in the November 2006 article cited above, IP spoofing is used to direct DNS response messages at third-party targets, including several Top Level Domain and Root name servers.

Attackers commonly use IP spoofing to impersonate "trusted" hosts as the first step in compromising a weakly secured host or server. Once compromised, these hosts can be used for any malicious purpose; one of the most common is to turn the host into a spam bot. Source identity spoofing can also be used in deception-based attacks against wireless access points, firewalls, VOIP/SIP proxies and other systems that filter traffic based on a value that can be easily forged (e.g., MAC address, hostname, email address).

■ **Lax System and Network Administration**—The SANS Institute identifies the most frequently-targeted OS, application and network components in its annual Top 20 list—but you could easily compile a list of common configuration errors that is twice as long and causes more security incidents. Unfortunately, this situation has not changed in the past four years. Many organizations still do not patch software in a disciplined manner; thus many incidents still exploit vulnerabilities for which fixes exist. Default settings in popular OSs and applications still accommodate ease of use at the expense of security, and far too many systems—including many wired and wireless network devices—operate with defaults.

Although more computers have antivirus and anti-spyware installed, the signatures are still not updated regularly. Too many file shares are still allowing guest access. Finally, many organizations still do not filter or analyze their outgoing (egress) traffic, affording attackers a vast freedom to exploit trojan back-channels.

■ **Sloppy programming is rampant**—Programming in the new millennium is as haphazardly practiced and oblivious to infection as surgery was during the military campaigns of the 18th and early 19th centuries. Just as no one needed a license in those days to practice medicine, no one needs a license today to develop software. Application and programming languages keep getting easier to use, resulting in more software development by inexperienced programmers who do not appreciate or practice secure coding techniques.

Inexperienced developers also are likely to download and use publicly shared code without questioning its origins or verifying what it does beyond what it claims to do. Understandably, many attackers now target poorly-written Web applications: this code tends to be relatively easy to crack and may be limited to a handful of Web servers, allowing exploitation to continue undetected for some time.

■ **New features trump security**—We still innovate with an "add security later" mentality. To corroborate this, we needn't look any further than voice over IP. In both public and enterprise networks, VOIP is commonly deployed with a weak authentication method that is susceptible to credential replay and spoofing attacks. VOIP's signaling and audio streaming protocols are also susceptible to numerous denial of service, interception and replay attacks. And numerous VOIP handsets harbor easily exploited vulnerabilities, from undocumented open maintenance ports to unencrypted management interfaces that have fixed or trivial passwords.

These examples illustrate that we are still unwilling to invest the time and expense to mitigate threats, including those for which we've developed solutions. These are symptoms of the same chronic problem we described in 2003: too many of us are still happy to mask rather than mitigate security shortcomings, burying our heads in the sand, and hoping that we are not the most appealing target when attackers look to do damage.

Hopeful Signs: Improved Systems And Software

On a brighter note, we do see some signs of progress. For example, many new business computers incorporate Trusted Platform Modules (TPMs), microcontrollers that store and protect keys and digital certificates against software attack and physical theft. As these platforms are more widely deployed, practically any Internet application will be able to leverage them. To do so, however, will require new identity management processes to issue the harder-to-spoof credentials which TPM can safely store. (For more about TPMs and other trusted computing developments see www.trustedcomputinggroup.org/news/Industry_Data/ESG_White_Paper.pdf)

Another hopeful sign is that operating system vendors are bowing to consumer pressure, and working to improve OS security. By default, the Mac OS X disables user access to core system functions and network access to all ports. Under the covers, OS X is based on a secure BSD

development environment that is continually strengthened through broad scrutiny by the open source community. With Vista and Longhorn, Microsoft has revamped Windows user account controls, hardened system services, added outbound filters, and incorporated anti-spyware. Even handheld device OSs like Symbian and Windows Mobile have grown new security knobs. And every major OS now offers automated updates. But the fact that flaws are still routinely patched in these newest OSs shows that much work remains.

The increased pressure to design, develop and publish more secure code extends beyond the OS suppliers and also has influenced commercial software developers and the open source community. For example, Oracle is one of many companies that encourage their developers to take ownership of coding flaws and to aggressively "hack" their own products. A growing number of startups are providing automated source code analysis as a commercial service.

In the open source world, broad code review and testing is increasingly being used to find and fix bugs before they can be exploited. For example, the U.S. Department of Homeland Security is underwriting a project that will study and refine techniques for identifying code flaws in Linux OS, the BIND name server, and the Firefox browser.

Network Security Improvements

Product support for secure network protocols also has also expanded. Since 2003, for example, the flawed Wired Equivalent Privacy protocol has been replaced with AES/802.1X-based Wi-Fi Protected Access v2. Solutions for securing offsite wireless include browser-based SSL VPNs and mobile VPNs that can survive roaming between networks and coverage holes. Unfortunately, even though these more secure protocols exist, adoption is slowed by pragmatic constraints like legacy system support and device management complexities.

Network security products also have evolved since 2003 to battle today's blended threats. Many single-function firewalls morphed into unified threat management (UTM) appliances that are able to apply intrusion prevention and antivirus to both inbound and outbound traffic at the network edge. Some UTM platforms also provide URL filtering, anti-spam and anti-spyware. This all-in-one approach is especially attractive to SMBs who lack the budget and staff to deploy more complex measures.

Increasing scrutiny at the firewall can reduce a network's risk exposure, but here again we find the gap between potential and practice. Deploying UTMs without risk assessment, policy development, and ongoing threat analysis can give companies a false sense of security. For example, many still use these more capable UTM appliances with default policies that filter only inbound traffic. And larger enterprises may always require point measures beyond UTM, such as high-volume mail firewalls that use source reputation and fingerprinting to fight botnets and image spam.

Another reason for optimism is Network Access Control (NAC), embodied by Cisco Network Admission Control, Microsoft Network Access Protection, and the Trusted Computing Group's Trusted Network Connect. These evolving architectures revamp the way we control access to corporate networks.

Instead of granting TCP/IP access to authenticated devices, NAC authorizes access based on verified user identity, the security state of that user's device, and policies that determine who may use which corporate resources, under what conditions. Furthermore, instead of simply blocking bad packets, NAC can automate remediation of infected or unpatched clients.

In theory, NAC can be applied to many different situations, from teleworkers with broadband PCs to wireless guests in conference rooms to employees who carry VOIP phones. In practice, each poses different challenges and goals. NAC adoption is in its infancy, as capabilities are gradually incorporated into the switches, routers, and firewalls that must enforce NAC decisions.

NAC adoption is in its infancy, as capabilities are gradually added to firewalls and switch/routers

Regulatory Impacts On Security

Since 2003, the pressure to comply with regulations like HIPAA, GLBA, SOX and PCI has frustrated many large IT shops, but these regulations also have contributed to overall network security. The financial and legal repercussions of non-compliance have prompted thousands of companies to track and control data access to a far greater degree than any prior imperative. More organizations now undergo vulnerability assessments to find and fix policy deviations before outside auditors spot them.

In short, regulations have raised "C-level" awareness and sensitivity to security threats and associated business risks. However, companies must avoid becoming overly focused on checklist compliance reports and passing audits—Don't stare so intently at one bit of bark on the nearest tree that you miss the larger threats devouring the surrounding forest.

IT professionals with security responsibilities would be wise to look for the hidden blessings in privacy regulations. Organizations may not be happy about being forced to implement privacy controls and audits, but these regulations certainly have raised security awareness. Seize this opportunity to improve your organization's security posture by assessing your vulnerabilities—not just in areas that require compliance, but across the board. Then use C-level interest in compliance to fund a comprehensive plan to mitigate those risks.

What Can We Do To Improve The Situation?

We have made some progress since 2003, although considerable work remains. Here are

Insist that your ISP provide secure routing, and that they block spoofed packets

some specific industry initiatives that we believe could mitigate the underlying weaknesses mentioned in this article:

■ **Deploy DNSSEC**—Harden your enterprise name service against attacks. The November 2006 *BCR* article on DNS Security summarizes a common deployment strategy. DNSSEC testbeds are already up and running in the UK, Sweden, Russia, Mexico, Bulgaria and many other countries, as are several generic top-level-domains, including .ORG. Participate in these initiatives and encourage your service providers and domain registries to deploy DNSSEC.

■ **Demand secure routing**—IETF committees are developing security enhancements for the OSPF and BGP routing protocols. The IETF OPSEC committee has done outstanding work defining a framework for operational security and a security best practices document. Ask your router vendor whether they are participating and when these enhancements will be available. Ask your ISP if they are adopting security best practices; better yet, make adoption a condition of your service agreement. (For more on security best practices in large IP networks, see *BCR*, March 2007, pp. 54–57.)

■ **Demand secure defaults**—Many organizations invest time and talent in developing secure client profiles, and even more in hardening servers. You are inviting attacks if you deploy and run systems with their out-of-the-box defaults. It would be better to insist that Microsoft, Apple, Red Hat and others provide an installation path that defaults to secure settings. Then we could start with locked-down systems and relax permissions only where our security policy dictates. Market pressures on vendors can prompt change.

■ **Prevent spoofing**—Most firewalls can enforce egress IP spoofing rules, yet few are configured to do so. SANS, RIPE and ICANN's SSAC recommend that you block inbound and outbound spoofed IP packets at your Internet firewalls and verify that your ISP is doing the same. Insist that spoofed packets be blocked as a condition of your service agreement. Moreover, avoid relying exclusively upon unverified source IP addresses (and other unverified source identities), by combining more secure network protocols with stronger authentication and more granular access controls.

■ **Strengthen authentication**—Look for ways to improve authentication, and offer only limited access for weakly-authenticated users. Take advantage of new laptop features like TPM chips and fingerprint readers. If you must use passwords with legacy systems, be sure to apply stronger authentication (e.g., tokens, smart cards, biometrics) on the front-end, and use single-sign-on. Home users should store text passwords safely in encrypted files (commonly called password safes).

■ **Narrow network access**—Improved network admission controls are long overdue. Don't use the competition among major vendors or the complexity of NAC deployments as an excuse to maintain the status quo. Incorporate NAC approaches where practical today—for example, many SSL VPNs can apply more granular user/group access controls, without requiring massive network NAC upgrades.

■ **Build security into all new Internet deployments**—Don't roll out new technology without security testing, policy development and risk mitigation. For example, don't distribute smart phones without a plan to track, secure and monitor their business content and usage. Don't turn up a wireless LAN without the ability to monitor wireless activity and stop unauthorized use. It is almost always less expensive to prevent a security incident than it is to deal with the aftermath.

Conclusion

Innovations in authentication, OS, software and network protocols show that we have the means to build a more secure Internet. Privacy and data security regulations give added impetus to enterprise security efforts.

On the other hand, we seem to be making the most progress primarily in the areas that are the easiest to solve, such as spam and anti-*X* filtering, while continuing to ignore the tougher multifaceted issues and the root causes which continue to undermine our systems, data and networks.

Four years hence, we hope that the state of Internet security will have improved. Ultimately, improving Internet security is a global problem, like curtailing greenhouse gas emissions. Unilateral improvements—in greenhouse gas emissions or in security—are simply not enough. We must all do our part and contribute where we can□

Companies Mentioned In This Article

Apple (www.apple.com)
Chase (www.chase.com)
Cisco (www.cisco.com)
Computer Emergency Response Team (www.cert.org)
Computer Security Institute (www.gocsi.com)
eBay (www.ebay.com)
Federal Bureau of Investigation (www.fbi.gov)
Internet Corporation for Assigned Names and Numbers (ICANN) (www.icann.org)
Infosecurity Europe (www.infosec.co.uk)
Microsoft (www.microsoft.com)
Oracle (www.oracle.com)
PayPal (www.paypal.com)
Red Hat (www.redhat.com)
Réseaux IP Européens (RIPE) (www.ripe.net)
SANS Institute (www.sans.org)
Symbian (www.symbian.com)
Trusted Computing Group (www.trustedcomputinggroup.org)