Global Information Assurance Certification Paper

GIAC Enterprises

A Practical Implementation of Defense In Depth and Concomitant Security Management Program

GIAC Security Leadership Certificate (GSLC)

Practical Assignment
Version 1.0

Dar-Ning Kung

Submitted March 7, 2003

Abstract

In an organization connected to the Internet for business operations such as e-commerce, both the security staff and the network administrators constantly face tremendous challenges from dynamic digital attacks upon the organization's network infrastructure, servers, workstations, and business services. For instance, the recent *SQL Slammer* worm infected thousands of systems around the world in a very short time, slowing down the Internet and disabling some bank ATMs. A combination of well-designed network security infrastructure features and security practices is essential to curtail future malicious attacks.

Since no single defense mechanism can defend against all types of threats and attacks, "defense in depth" is the best solution for safeguarding an organization's IT systems and digital assets. This paper presents a network security architecture including routers, switches, and firewalls to exemplify the "defense in depth" concept. The discussion shows how the internal network is separated from the Internet and how various additional protective measures are employed to counter attacks from cyberspace. In addition, the paper discusses sound, practical security management practices for both network and host protection and demonstrates how they significantly enhance the organization's overall security posture. Finally, as a case study in organizational security, it is shown that a combination of sound security practices with judicious changes to the IT infrastructure can successfully defend against both known and unknown malicious agents.

Introduction

Widespread and ever-growing Internet usage has motivated companies to expand their business operations into e-commerce at a breakneck pace in order to catch the wave of improved efficiency and effectiveness and facilitated customer access, all hopefully leading to increased revenue. The desire to increase exposure and reach out to new customers must be balanced with the need to provide safe and secure access to in-house staff and protect resources from competitors, dedicated attackers, and even innocent interlopers. This report begins by describing the nature of attacks from cyberspace, which is often compared to the "wild West" of the 1880s, and the business impact that can result from these attacks. A network security infrastructure is proposed to protect both the network and the hosts while carefully balancing the bewildering state of the art with ease of integration and management. Without presenting exhaustive details about available and emergent technologies, this paper demonstrates that the combination of "defense in depth" with the successful implementation of sound security practices can satisfactorily protect an organization's assets at reasonable cost. It is then demonstrated how this two-pronged approach successfully defended against the recent *SQL Slammer* worm in particular and wider classes of attacks in general.

Notably, though the business requirements for disparate organizations – such as government agencies, commercial firms, and financial institutes – may be quite different, the information security objectives of confidentiality, integrity, and availability (CIA) are the same for everyone. A fortunate result of this fact is that an acceptable secure network design can typically be based directly upon the existing network infrastructure. Yet, the infrastructure must be capable of growing and evolving to accommodate changing business needs, emerging security technologies, acceptable risk threshold, and applicable legal, financial, and scheduling constraints. The evolving secured network combined with a well-thought security program can achieve the objective of organization's information security—protecting the host operating systems, applications, and data against attacks and thereby ensuring information confidentiality, integrity, and availability per applicable policy and law.

Selection criteria: the business threat impact

The CERT Coordination Center publishes security statistics to bring public attention to the threats encountered in cyberspace and the dramatic upward trend thereof. For instance, though fewer than four thousand incidents were reported in 1998, more than eighty thousand were reported in 2002, while the number of published vulnerabilities increased from approximately three hundred in 1998 to in excess of four thousand in 2002. These numbers scarcely begin to demonstrate the dramatic increase of exploitable threats from which an organization's information assets may suffer. Consequently, all organizations must resign themselves to the reality and imminence of these threats and respond by building secure networks, fortifying in-house security policies, and deploying available technology to vigorously defend themselves against potential perpetrators of cyber crime.

Threats to organizational IT assets stem from a variety of sources. Internal users, who already possess authorized access to the internal network, might include disgruntled employees, corporate spies, or merely clumsy computer users whose inadvertent mistakes may be damaging or even crippling. A proliferating array of threats originate from the exposure of an organization's public host machines to the global Internet – machines that can be probed and explored by hackers as well as legitimate customers, both current and potential. Finally, dial-in and wireless services can perhaps be exploited to gain access to internal resources in as yet unknown ways.

In particular, it is crucial for IT mangers to understand hostile threats and types of attacks from cyberspace and ask themselves whether the organizations can afford the business impacts and financial losses caused by successful attacks. For instance, an attack may cause an organization services interruptions lasting from several hours to many days due to

compromised machines or infected hosts. If sensitive or financial information has been compromised, the company may be required to invest a great deal of time and money to fix the problem. In addition, the aftermath may include losses of millions of dollars' worth of revenue, exposure of confidential business secrets, encumbrance of public relations, and loss of customer confidence. The organization may even find itself the target of legitimate lawsuits by customers seeking monetary or punitive damage for mishandling of their personal or proprietary corporate data that had being entrusted to an organization as a 'safeguarded' business entity with well-established security-enabled business operations.

Network security infrastructure: design and analysis

Much has been written about various possible security architectures and the tradeoffs among them. Despite the considerable variety of approach and the fact that architectures must obviously be tailored to the needs of the specific environment, a three-zone architecture seems to offer a good balance between security and complexity. The three zones that comprise the security architecture are the perimeter zone, the demilitarized zone (DMZ), and the private (internal) network zone. An implicit trust relationship exists whereby increasing trustworthiness is ascribed to both hardware and users as one navigates inward from the outermost perimeter to the inner sanctum where corporate secrets and sensitive customer data are housed. An important feature of this configuration is that, in order for traffic from the Internet to enter the organization's private network, it must traverse two firewalls. Typically, the firewalls are built by two different vendors so that—postulating an easily exploited vulnerability in one of the firewalls—the probability that the second firewall can effortlessly be penetrated in precisely the same manner is dramatically reduced.

Figure 1 shows the proposed network security infrastructure. The layered approach has the following advantages: access to external services at DMZ does not need affect the private network that is protected by a firewall; public services are isolated in the DMZ; and internal network has two layers of firewall protection. Also, the externally accessible servers in DMZ prevents a compromised server from analyzing the traffic from/to the internal network and deterred further damage to the internal network. Note that, even though the proposed architecture provides a high degree of network security, hosts in the internal network are still require harden secured configurations to minimize the impact if firewalls are compromised. This issue is discussed in more detail in an ensuing section of this report.

Figure 1. Proposed Network Security Architecture

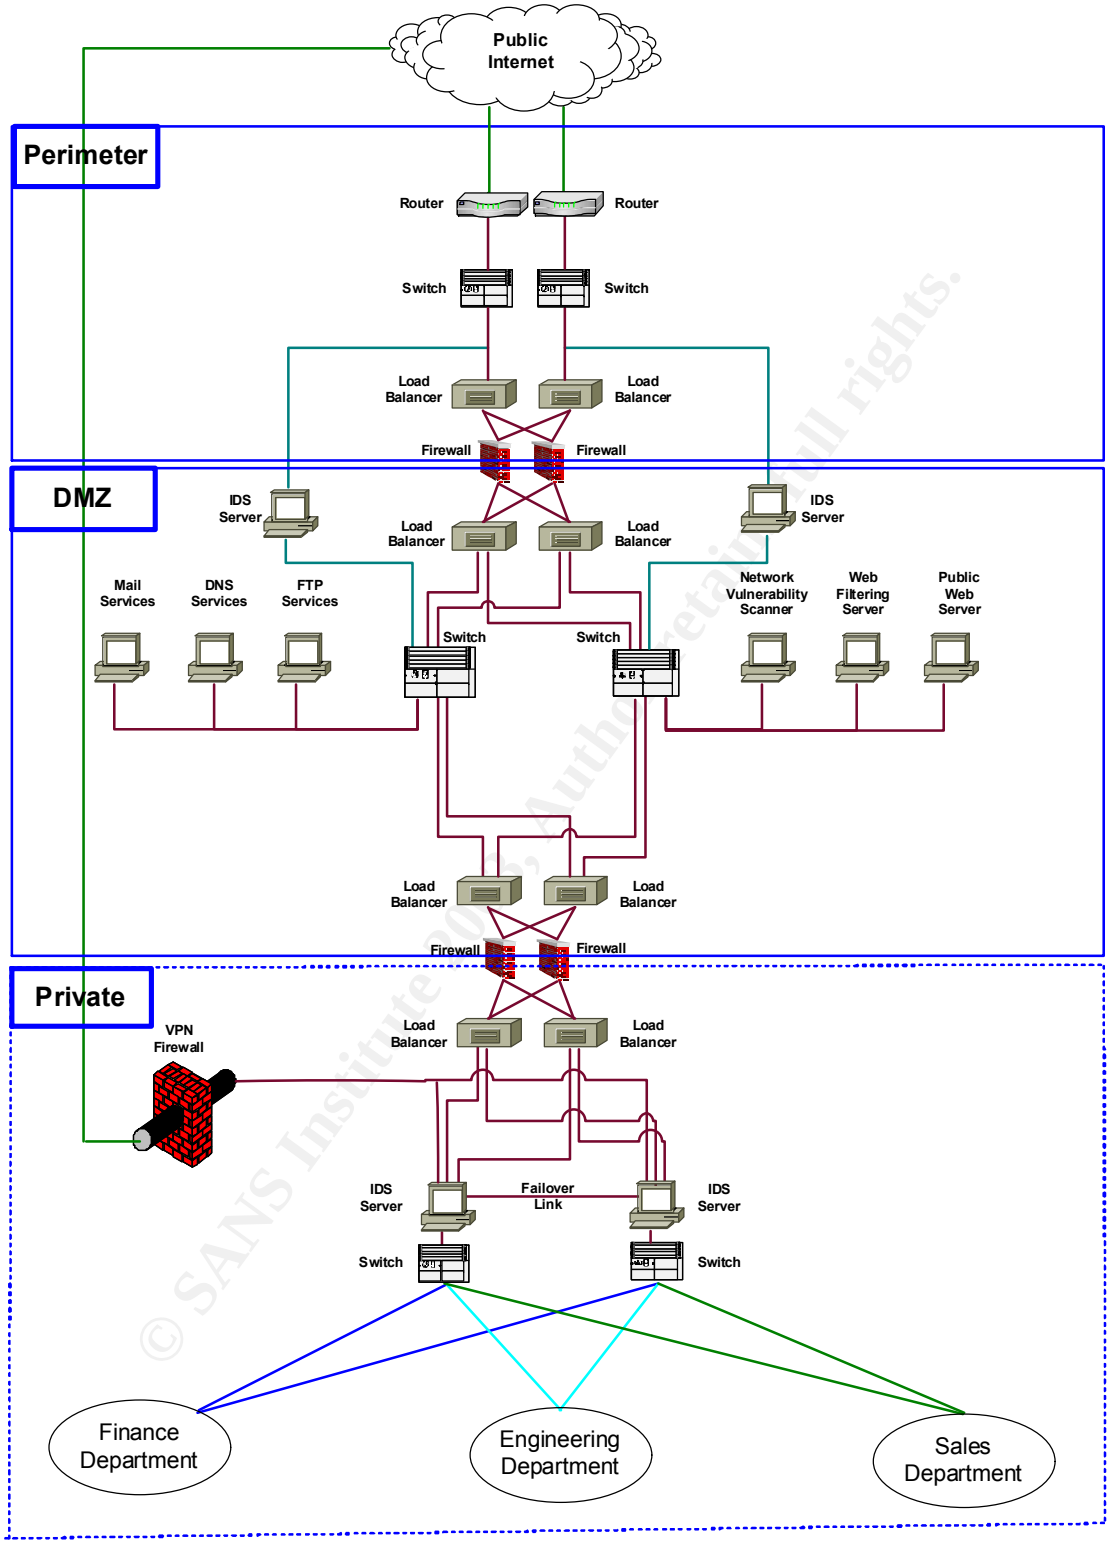

The components of the network security architecture and their individual zone memberships are characterized in Table 1.

Table 1. Network security architecture components

Zone	Device	Function
Perimeter	Router	Determines whether a packet inbound from the Internet should be accepted and forwarded to the next zone
	Switch	Provides full connectivity and separates logical and physical network broadcast domains using VLAN technology
	Intrusion detection system (IDS)	Signature-based monitoring of known and unknown penetration exploits
	Load balancer	Ensures maximum uptime through guaranteed public connection redundancy
	Cisco *PIX* firewall	Separates the untrusted world from the demilitarized zone
DMZ	Load balancer	Ensures maximum uptime through guaranteed public connection redundancy
	Switch	Provides full connectivity
	FTP Server	Provides content-based filtering of e-mail attachments, including viruses and worms
	Mail Relay	Provides additional level of defense against mail-based attacks, *e.g.*, spamming
	Public DNS server	Provides restricted public access to organization's internal structure and naming hierarchy
	Public Web server	Provides restricted public access to organization's hypertext-based information resources
	Load balancer	Ensures maximum uptime through guaranteed public connection redundancy
	NetScreen firewall	Separates the demilitarized zone from the private network
Internal network	Database server	Provide organization either internal or external transactional database operations

Zone	Device	Function
	Internal application server	Provide organization either internal or external special applications such as Customer Relation Management, Payroll system, and HR system
	"Extranet" VPN	Provide secure, encrypted connections between an organization's private network and remote users or office through Internet
	"Extranet" firewall	Provide filtering, blocking and remote Virtual Private Networking access for an organization's private network
	Private DNS server	Provides full access to organization's internal structure and naming hierarchy
	Private Web server	Provides full access to organization's hypertext-based information resources

Security is enforced throughout the infrastructure, thereby providing defense in depth. The multi-layered defense will ensure that a hacker must expend a great deal of effort to penetrate the defensive mechanisms and thereby initiate compromise of a protected host. In addition, if a hacker must take more time to probe and potentially identify weaknesses of the network or hosts, the probability that intrusion detection system (IDS) alarms will be triggered is increased substantially. Such strength of mechanism serves as a psychological deterrent as well, inducing the hacker to seek a weaker victim.

The network security architecture emphasizes reliable response to inbound HTTP access over response to outbound traffic. Also, beyond the traditional focus upon confidentiality and integrity, the architecture specifically addresses the availability issue by incorporating judiciously applied fail-over mechanisms that defray single points of failure. Specifically, the two pairs of load balancers on either side of the outermost firewall ensure that critical network services are available in the event of all but catastrophic universal failure. The load-balancing infrastructure not only improves the overall performance on network traffic, but also establishes high scalability for an enterprise network infrastructure, thereby accommodating future growth in response to evolving business requirements.

The configuration details of several of the foregoing network components are explained in Table 2.

Table 2. Secure network component configuration details

Component	Configuration details
Perimeter router	Denies ICMP redirects, unroutable private address access, and spoofed addresses; *Tripwire* installed to monitor changes to configuration files; integrity checks run as needed
DMZ router	Permits only SMTP, DNS, HTTP, and HTTPS traffic destined for

Component	Configuration details
	mail, naming, Web, and secure Web services, respectively; *Tripwire* installed to monitor changes to configuration files; integrity checks run as needed
Private router	Allows encrypted VPN traffic from remote access firewall only and permits access to only specific portions of the protected network; *Tripwire* installed to monitor changes to configuration files; integrity checks run as needed
Cisco *PIX* firewall between perimeter and DMZ	Allows only management traffic to firewall; allows only HTTP, HTTPS, DNS, VPN, and application-specific access to specific IP addresses; drops all other traffic
NetScreen firewall between DMZ and private network	Allows only management traffic to firewall (including traffic from outside monitoring firm); allows VPN access to specific IP addresses; drops all other traffic
IntruVert IntruShield intrusion detection system (IDS)	Configures the IDSs in-line mode in the private zone to drop or block a single packet or a single session between the attack source and destination Configures the IDSs standard mode in the DMZ to monitor the suspicious traffic before and after the firewall for event correlation

Certain ingredients of the network security architecture are of sufficient complexity that they merit specific investigation above and beyond the foregoing configuration schedule. The first of these is the virtual private network (VPN) service provided by both firewalls. By default, the VPN maintains a table of registered media access layer (MAC) addresses for which laptops have been issued. Nortel group password token authentication is relied upon to verify each user's claimed identity. This authentication scheme is part of the standard issue for all mobile users within the organization and relies upon one-way encryption of a fixed password string.

The second component that deserves more focused attention is the IntruVert *IntruShield* intrusion detection system (IDS). The IDS provides real-time, gigabit-speed network intrusion detection and prevention solution and subscribes to automatic signature updates in real time. The IDS sensor monitors traffic across all three zones, providing real-time detection and prevention mechanisms that can drop a single packet or session; initiate TCP reset or ICMP unreachable response packets; automatically reconfigure the firewalls to incorporate dynamically generated rules that thwart likely attackers; capture and log packets both before and after a detected attack for detailed post-incident forensic analysis; and send a real-time alert to the IDS manager via e-mail or pager.

Finally, in order to rein in the personnel and economic costs associated with monitoring all security-related devices, including log post-analysis and event correlation, an enterprise-wide security event management software package, *e-Security*, is installed. The software aggregates, standardizes, analyzes, and reports security event information from security devices organization-wide to a centralized console in real time.

Security management requirements

The network security infrastructure can be compared to a skeleton: while it provides a necessary framework for the organism, it must be fleshed out, so to speak, in order to support full functionality. The flesh of the organic structure takes the form of comprehensive security management. Security management begins with proper configuration of hardware and software by system administrators. Unused services must be disabled; access to critical system directories and files must be secured; and known system vulnerabilities must be systematically patched or otherwise eliminated. Administrators must also establish a standard build with properly secured configuration that meets the organization's business needs as a baseline system.

A particularly good practice is to unify system configurations with standard builds for hosts connecting to the organization's private network, thereby supporting group policies by which systems can be automatically brought into conformance en masse with organizational standards. For instance, it is possible to apply a group policy to a large quantity of Windows 2000 desktops simultaneously when an organization established accounts in Active Directory. Software Update Server (SUS) downloads hot fixes and patches directly form Microsoft's web servers. System administrators can perform testing on newly downloaded hot fixes and patches. After the testing, system administrators can then apply a group policy and configure the Windows workstations for the Automatic Update Service within the organization. The benefits for performing group policies are the followings:

- Service packs, hot fixes, and patches are being applied to desktops, which greatly reduces exposure to known security vulnerabilities.
- Workstation stability can be improved.
- A much more uniform "leveling" of desktops to current security standards can be enforced and maintained.
- Greatly reduced administrative overhead can be achieved.

Importantly, any newly deployed system must be subjected to a vulnerability scan to ensure that the system is brought into the network without any holes that can either compromise that host or, through exploitation of transitive trust relationships, compromise wider sets of resources. Indeed, scanning of individual hosts on a periodic basis can be viewed as the linchpin of host-oriented network security, much as proper configuration of firewalls is the linchpin of protocol-oriented network security.

Table 3 discusses representative of protocol-centric and host-centric tests that constitute the organization's ongoing vulnerability assessment and that, ideally, are conducted from both outside and inside the organization's private network.

Table 3. Host and network scanning program elements

Program element	Description and purpose
Network mapping	Identifies all workstations, servers, network and communication devices as well as unauthorized machines connecting to the organization's network. In addition, the open ports and unauthorized services can also be identified. Any discovered deficiencies are required to be corrected in a timely manner.
Network-based vulnerability scanning	Identifies all hosts, workstations, and network devices' vulnerabilities. Most scanners also provide recommendations on how to mitigate discovered vulnerabilities. Because newly discovered system and application come out everyday, the vulnerability scanning has to be performed regularly, and each scan must include the signatures of newly identified vulnerabilities. The network scan focuses on the advertised network services and vulnerabilities of each box. However, it may not be able to identify the specific OS and the associated latest patch for the box. Therefore, the network scan results may contain many false positives.
Penetration testing	Determines how vulnerable an organization's network can be is to penetration by a skilled intruder and how much damage can result from such an intrusion. This test provides an objective view of whether the implemented security infrastructure can achieve the security requirements articulated in the security policy and thereby effectively support secure business operations.
Host-based scanning	Determines what host vulnerabilities are exploitable via the network. It is a complement to network scanning. The host scan tool must be individually installed on each machine, and a different version of the scan tool may be required for different OS. In addition, the host scan has to be performed at the same machine, not from another network machine, so the process requires an administrator login to perform security checks on the host.
Host-based intrusion detection	Detects unauthorized file access or modifications. The process examines the log files to identify potential host misuses.
Virus detection	Detects and deletes viruses from inside or outside of the internal network. Virus detections also can use a layered approach, e.g., *VirusWall* at the perimeter to scan inbound or outbound e-mail messages and their attachments, and classical anti-virus software at the mail server and all workstations. The key to successful virus detection and prevention is vigilant maintenance of up-to-date viral signature databases. If possible, the signature updates can be an automated real-time process.
Password cracking	Verifies whether users' passwords are easy to guess, thereby enforcing strong password use and enhancing password policies. Practicing good password procedures is a key defensive measure in protecting valuable organization information resources.
War dialing	Detects unauthorized modems and prevents unauthorized access to the internal, private network. System administrators can use commercial war dialers, also known as modem scanners, to

Program element	Description and purpose
	identify unauthorized modems on an enterprise network. Such modems can provide an intruder easy access to an organization's intranet.

Even if an organization has established a strong security infrastructure, there is still a risk that a security incident may occur. For instance, if an administrator observes suspicious activities logged on a server or on a private network, he/she needs to evaluate and analyze the problem and identify whether a host has been compromised. If a compromise is identified, the incident response team may be called for further investigation such as forensic analysis and damage assessment. After an incidence analysis, a security administrator should document the findings and the lessons learned thence and share this knowledge with other security, network, and system staff throughout the organization. At the conclusion of the exercise, a revision to organizational policies may be in order to prevent either a reoccurrence or a similar incident in the future.

The foregoing mechanisms, equipment, and policies are, in actuality, only a part of the organizational approach to repeatable, documented security that begins with policies and procedures to ensure that good security practices are a part of each person's regimen. Organizational policies clearly define what is allowed and what is not allowed. A policy covering an entire organization may consist of many documents, each specifying the scope, ownership, roles, and responsibilities of a certain segment of the user or administrator population. It also may include guidelines on the management of expected or unexpected risks. In general, the security officer's or manager's duties can be said to consist largely of prioritization and implementation of the documented security program based upon organizational policies and resources, not only ensuring satisfactory adherence to what is current, but also keeping an eye on future improvement of the organization's security posture.

Applicability of security infrastructure: the *SQL Slammer* worm

On 25 January 2003, the *SQL Slammer* worm, the fastest known computer epidemic in history, spread throughout the Internet and infected more than 90 percent of vulnerable hosts within ten minutes. The worm targeted a known vulnerability in Microsoft *SQL Server 2000*, automatically searching for vulnerable systems and infecting them via those services and presently resulting in an epidemic distributed denial of service (DDoS) attack that completely saturated available network bandwidth. Given the remarkable speed at which the worm spread, it was simply impossible for security administrators of vulnerable networks and hosts to identify the problem and implement remedial measures.

The proposed network security infrastructure, judiciously operated and managed with sound security practices, established a robust three-layer defense against the worm:

- Access control on the edge router that has a point-of-presence (POP) to the Internet) and access control list (ACL) rejected packets that were destined to unused host addresses inside the organizational subnets from unknown or unpredictable addresses.

- Packet filtering logic on the firewalls that blocked TCP ports 1433 and 1434 (the ports used by the *SQL Slammer* worm) prevented remote users from directly accessing any organizational SQL servers from the Internet. Had an administrator noted a legitimate business requirement for direct outside access to the SQL server, a virtual private network (VPN) solution would have provided adequate secure access to accommodate the administrator's needs.

- A periodic scan on all organization subnets gave administrators an opportunity to identify systems' vulnerabilities and request system owners to apply patches remedy them. The known vulnerability associated with this *SQL Slammer* worm should have been completely remedied insofar as Microsoft had announced the vulnerability in July 2002 and had provided a patch to accompany the alert bulletin. Consequently, even if an *SQL Slammer* worm had entered the internal network, there would be no vulnerability for the worm to exploit.

Conclusions

All organizations connected to the Internet face increasing risks and threats originating from cyberspace as well as inside of the internal network. If an organization does not diligently enforce a holistic information security policy, it will only be a matter of time before an organization's network is compromised, and the potential damage caused by the security incident may be tremendous.

The proposed network security infrastructure is a solution that addresses the functional and business requirements for a medium to large enterprise network. It is not meant to be a best infrastructure for all business. Rather, the proposed infrastructure adapts the defense-in-depth concept to securing a typical organizational configuration. The security infrastructure and management program will, in general, vary depending on business requirements, cost and resource constraints.

Finally, the proposed network security infrastructure and accompanying security management program require a commitment not only from the executive staff and IT managers of an organization, but strong engineering, administrative, and even end user support . Everyone in the organization must be acutely aware of

the importance of network and host security and of the criticality of his own role in its enforcement. Only through complete understanding of roles and responsibilities and eternal vigilance can an organization's IT security posture be sufficient to withstand the attacks of unknown scope and complexity that wait just over the horizon.

References

[1] Allen, Julia H., The CERT Guide to Systems and Network Security Practices. Boston: Addison-Wesley, 2001.

[2] CERT Coordination Center, "CERT/CC Statistics 1988-2002", URL: http://www.cert.org/stats/cert_stats.html, (March 5, 2003).

[3] SANS Institute, Course Material from "Track 12 – Security Leadship for Managers", 2002.

[4] SANS Reading Room Documents, URL: http://rr.sans.org/index.php, (March 5, 2003).

[5] Wack, J., Cutler, K., and Pole, J. Guidelines on Firewalls and Firewall Policy, Special Publication SP 800-41, National Institute of Standards and Technology, January 2002.

[6] Wack, J., and Tracey, M. Guidelines on Network Security Testing, Special Publication 800-42 DRAFT, National Institute of Standards and Technology, February 2002.

[7] Weaver, Nicholas C., *et al.*, "The Spread of the Sapphire/Slammer Worm", URL: http://www.cs.berkeley.edu/~nweaver/sapphire/, (March 5, 2003).