IAnewsletter

The Newsletter for Information Assurance Technology Professionals

Volume 7 Number 3 • Winter 2004/2005

The Cyber Conflict Studies Association

also inside—

- Preventing Widespread Malicious Code
- The Future of Network Intrusion Detection
- IPv6—The Next Generation Internet Protocol
- The Importance of High Quality IA Metrics
- DEFCON 12 Security Conference
- Evidence-based Health Care and IA

contents

feature

IA initiatives

in every issue

About IATAC & the *IAnewsletter*—

IAnewsletter is published quarterly by the Information Assurance Technology Analysis Center (IATAC). IATAC is a Department of Defense (DoD) sponsored Information Analysis Center, administratively managed by the Defense Technical Information Center (DTIC), Director, Defense Research and Engineering (DDR&E).

Contents of the *IAnewsletter* are not necessarily the official views of or endorsed by the U.S. Government, DoD, or DDR&E. The mention of commercial products and/or services does not imply endorsement by the DoD or DDR&E.

Inquiries about IATAC capabilities, products, and services may be addressed to—

IATAC Director:	Gene Tyler
Acting Deputy Director:	Jim Peña
Inquiry Services:	Peggy O'Connor

IAnewsletter Staff—

Creative Director:	Christina P. McNemar
Art Director:	Ahnie Senft
Designers:	Ricardo Real
	Dustin Hurt
	Kathy Everett
Editorial Board:	Jim Peña
	Ronald Ritchey
	Tara Shea
	Gene Tyler

IAnewsletter Article Submissions

To submit your articles, notices, programs, or ideas for future issues, please visit http://iac.dtic.mil/iatac/IA_newsletter.html and download an "Article Instructions" packet.

IAnewsletter Address Changes/Additions/Deletions

To change, add, or delete your mailing or E-mail address (soft-copy receipt), please contact us at—

IATAC
Attn: Peggy O'Connor
3190 Fairview Park Drive
Falls Church, VA 22042

Phone: 703/289-5454
Fax: 703/289-5467

E-mail: iatac@dtic.mil
URL: http://iac.dtic.mil/iatac

Deadlines for future Issues—

Spring 2005	February 1, 2005

Cover design:	Ricardo Real
Newsletter design:	Ahnie Senft

Distribution Statement A:

Approved for public release; distribution is unlimited.

IATAC Chat

Gene Tyler, IATAC Director

Good day fellow cyber security professionals! As the newly minted IATAC Director, I want to take a few minutes and "thank all of you" who have warmly welcomed me and eagerly guided my development as I assume these new duties.

Acceptance as the new Director has been encouraging and humbling, and I must say I am a bit awe struck. I never really grasped the full impact of all the IATAC activities before my arrival. Three years prior to coming to IATAC, I served as the Director of the Defense-wide Information Assurance Program (DIAP) and concurrently completed the last of my 35 years of U.S. Army service. In my capacity with the DIAP, I served on the IATAC Steering Committee. The DIAP and IATAC have very similar missions, goals, and constituents—both serve to expand and develop the body of IA and IO knowledge. We can easily expand that thought to include all the IA professionals from industry, academia, and government, as all have the same end goal—secure and safe networks and computing environments.

Bob Lamb passed on the reins of a strong and vibrant organization. We are committed to further developing the IA and IO endeavors that relate to IA. Recently, the IATAC Steering Committee challenged us to accomplish a number of important tasks. I would like to highlight three here—

- Conduct an overlap and gap analysis of the DoD IA Strategy, IA Component to the GIG Architecture, the IA Hard Problems List, and the DoD CND Roadmap
- Enhance interactions with IA Centers of Academic Excellence and Military undergraduate and graduate institutions
- Provide IATAC Steering Committee the list of IATAC Subject Matter Experts (SME)

These three points have broad appeal to the "readership" of the *IAnewsletter*. I suspect you have heard of them, and it is my intent to highlight these issues from time to time in the *IAnewsletter*.

The first topic will involve a significant analysis of key issues in the IA world. The overlap and gap analysis work is extremely complex and will require a major effort to ensure the results are well coordinated and synchronized throughout the DoD IA community. Adding to the complexity, the Joint Staff J6 is developing an IA Campaign Plan—our gap analysis will consider this work too. Our intent is to map the results into a State-of-the-Art Report (SOAR) and present the findings as a part of the February 2005 IA Workshop in Philadelphia. Our timeline is aggressive and to meet it, we will prebrief the Steering Committee membership at our January 2005 meeting. We will be working closely with the OSD Staff, DIAP, Joint Staff J6, USSTRATCOM, DISA, and NSA.

The IA Centers of Academic Excellence and the SME databases are important efforts, as they serve as capacity and infrastructure building blocks—that is educate, develop, and capture IA professionals. There are over 60 Centers of Academic Excellence. Every DoD IA professional should be aware of them and the IA Scholarship Program. These institutions offer super opportunities and DoD is working to integrate your educational experiences with the institutions' programs to maximize learning and create opportunities for each of us. The NSA is DoD's Executive Agent for the DoD IA Centers of Academic Excellence, and they have a solid program that offers something for all. Go to http://www.nsa.gov/ia/academia/caeiae.cfm to find out more. Future issues of the *IAnewsletter* will contain a section highlighting selected institutions and relate their work to DoD activities.

Building a strong and vibrant IA professional community requires identifying the "graybeards." IATAC's SME database is leveraged when we receive a query from you. Often, we immediately know the answer to your question and we respond quickly. If your question is more complex, we discuss it with a SME, who guides our researcher in the proper direction. We are always looking to expand our list. You can learn more about the IATAC SME program by going to http://iac.dtic.mil/iatac/sme.html. Future editions of the *IAnewsletter* will highlight a selected SME. Our intent is to give credit where credit is due and our IA community SMEs serve an important role in developing and expanding the IA body of knowledge.

In closing, I receive lots of mail from you, the readers. Future editions will highlight some of your comments to me. So keep your comments and questions coming and we will do our best to respond. I want to ensure we maintain a balanced approach to serving the operational, personnel, and technical portions of the IA community and the readers of the *IAnewsletter*. I believe this effort will reinforce our commitment to stay plugged in to you—our customers. ■

The Cyber Conflict Studies Association

by Robert Gourley and John Casciano

Founded in 2003, the Cyber Conflict Studies Association (CCSA), is a Not-for-Profit, national membership organization devoted to the study of issues related to conflict in the Information Age. Our objectives can be summarized in three areas—

- First, to educate the public on the range of issues that touch on cyber conflict. The Association conducts a series of educational events including symposia, workshops, and luncheons throughout the year, and provides speakers and panelists to other organizations upon request.

- Second, to provide a forum for the debate of public policy related to the full range of cyber conflict issues—moral, ethical, legal, political, economic, sociologic, military, scientific, etc. The Association sometimes teams with other organizations, academia, and government to provide policy forums for the debate of relevant issues. CCSA also establishes independent task forces to examine particular public policy agendas.

- Third, to sponsor and/or perform scholarly research on conflict in the Information Age. The Association works with public and private entities to develop research topics, secure funding, and administer programs. A major thrust is to encourage cross-disciplinary research and dialogue.

Although CCSA eventually plans to publish an independent journal, our primary method of communicating our work today is our Web site at http://www.cyberconflict.org and our quarterly newsletter.

The CCSA has sponsored several conferences and symposia where we have tackled strategic issues such as the economic impact of cyber conflict and the strategic imperatives for new research.

In the seven years since the issuance of the President's Commission on Critical Infrastructure Protection (PCCIP) Report, much has changed in the external environment, in the threats, and in our national response. A new Administration, the tragic events of September 11, 2001, the Global War on Terrorism, the Patriot Act and other legislation, the creation of the Department of Homeland Security, the evolving technologies, and many other factors have all had an impact on the cybersecurity environment. New organizations, new players, and a new National Strategy to Secure Cyberspace (February 2003) are now driving the responses recommended in the PCCIP Report.

This is a propitious time to review and assess the impact of the PCCIP Report and examine and evaluate its impact. This symposium is designed to do just that by bringing together former and present players in the government, industry cybersecurity leaders, and academicians to discuss the report and the impact it has had.

The symposium is structured around some of the principal assumptions, findings, and recommendations of the PCCIP Report—

- To what extent were the postulated threats valid? What was missed or poorly understood? What has changed?

- What about the Commission's assumptions? Were all of them valid? Were any missed? Are there new ones driving our response today?

- How has the role of government evolved in responding to cybersecurity challenges? Is government taking the lead and showing the way?

- What have the various departments and agencies accomplished since 1997 in dealing with the threats and vulnerabilities?

- How has the public-private partnership envisioned in the Report worked? Is there adequate sharing between industry and government and among the Information Sharing and Analysis Centers (ISACs)?

- What are the legal and economic considerations attendant to this issue?

- How well have the government's education, outreach, and research and development programs responded to the PCCIP's recommendations?

All indications are that this conference will be watched by a wide range of very senior government officials. Our hope is to produce results of a quality that will have a positive impact on the future of cyber conflict. We know, however, that reaching this goal will only be possible by the continued engagement of those active today in the realm of cyber conflict, which includes the vast majority of *IAnewsletter* readers.

To learn more about the CCSA, to join or register for our conferences, or to volunteer your assistance please visit http://www.cyberconflict.org. ■

About the Authors

Robert Gourley

Bob Gourley is the Director of Strategy and Planning for Northrop Grumman's Intelligence Systems Division. A retired Naval Intelligence officer, Mr. Gourley was the first Director of Intelligence (J2) for the Joint Task Force for Computer Network Defense (JTF–CND, later JTF–CNO). He holds three masters degrees including an MS in Computer Science from JMU with an emphasis in Information Security. Mr. Gourley may be reached at bob.gourley@ngc.com.

John Casciano

John Casciano is a 32-year veteran of the United States Air Force, having retired in 1999 as a Major General. After retirement, he managed businesses focused on Information Warfare and Security for two large companies before starting his own firm. John has both a Bachelor's and Master's degree from Georgetown University and has attended several executive programs in National Security and Business at Harvard and the Wharton School. Mr. Casciano may be reached at Graystr@aol.com.

Preventing Widespread Malicious Code Infections

by Karen Kent

Malicious code has become one of the greatest threats against the information technology resources of most organizations. A single widespread malicious code incident can potentially cause serious damage to systems throughout an organization in a matter of minutes, requiring weeks of recovery efforts. This article describes the current malicious code threat environment, explains the need for a layered defense against malicious code, and identifies best practices for preventing the most common types of widespread infections, supported by requirements and recommendations from DISA's Security Technical Implementation Guides (STIG), and incident handling guidance from the National Institute of Standards and Technology (NIST).

The current malicious code threat environment

Nearly all recent major malicious code incidents—those that have significantly affected many organizations around the world—have involved worms. Worms are self-contained instances of malicious code (unlike viruses, which infect existing files). Two types of worms have been most prevalent in 2004—mass mailing and network service worms. As the name indicates, a mass mailing worm spreads through E-mail. Some mass mailing worms take advantage of E-mail client or operating system vulnerabilities, while others simply rely on unwitting users receiving and running the worm. Examples of recent mass mailing worms are *Beagle*, *Mydoom*, and *Netsky*. A network service worm typically spreads without any user assistance. It scans or connects to systems on a port associated with a particular service, identifies unsecured or otherwise vulnerable services, then takes advantage of the weaknesses to infect the hosts. *Sasser* and *Witty* are recent network service worms.

Several years ago, most worms were just nuisances, causing relatively little damage. For example, in May 2000 the *Love Letter* worm reportedly infected millions of systems, but system damage was limited to overwriting some graphics files and altering users' Internet Explorer start pages. *Love Letter*, which spread primarily through mass mailing, also

overwhelmed many E-mail servers and networks, but most organizations were able to stop the worm within hours. Today, most major mass mailing worms still cause a serious strain on E-mail infrastructures, and both mass mailing and network service worms often generate very high volumes of network traffic. During widespread infections, many organizations experience E-mail and/or network reliability issues for hours or days, sometimes even weeks.

Although the impact of a widespread malicious code infection on E-mail servers and networks has not fundamentally changed over the years, such infections have become far more frequent, and multiple new malicious code threats sometimes occur simultaneously. Also, many of today's worms can cause much more serious consequences to workstations and servers than overwriting graphics files and altering home pages. The following items illustrate the potential damage caused by recent worm infections—

- Creating backdoors that can be used by attackers and other malicious code to gain unauthorized remote administrator-level access to workstations and servers.
- Disabling security services such as antivirus software and personal firewalls, which can facilitate other successful attacks and unauthorized access.
- Participating in a distributed denial of service attack against another host or network.
- Causing systems (intentionally or inadvertently) to crash or shut down and restart repeatedly, making them difficult or impossible to use.

Even though backdoors and other system alterations caused by malicious code could potentially lead to additional damage and unauthorized access, organizations typically have sufficient time to isolate infected systems before attackers can further exploit or damage them. This was not possible with the *Witty* worm, released in March 2004. A network service worm, *Witty* targeted a vulnerable service used by certain intrusion detection and prevention software products, and attempted to overwrite data on the systems' hard drives, effectively wiping the systems out. *Witty* spread so quickly—in a matter of minutes—that most vulnerable systems were already affected before administrators had a chance to react. Although the number of systems directly affected by *Witty* was not that high because of the low prevalence of its targets, a similar worm could easily be constructed that would take advantage of a vulnerability in a popular operating system or desktop application, potentially wiping out the hard drives of hundreds of thousands (perhaps millions) of systems in a matter of minutes.

Malicious code can be spread through many mechanisms besides mass mailing and network service worms. For example, malicious Web sites could attempt to infect visiting hosts. Services such as instant messaging and peer-to-peer file sharing have also become a common way to transmit malicious code by tricking users into receiving malicious code through file transfers. However, these methods generally do not cause widespread infections to occur rapidly. Accordingly, this article focuses on prevention strategies for mass mailing and network service worms because they are the infection mechanisms most likely at this time to breach many hosts rapidly.

The need for multiple infection prevention measures

Because malicious code threats are becoming more damaging, it is more important than ever to prevent malicious code incidents whenever possible. The primary prevention mechanism is antivirus software; however, although it is often effective, it cannot stop all infection attempts. Antivirus software is excellent at identifying known threats, but often not so good at recognizing new threats. Most antivirus signatures are based on the characteristics of known malicious code—they may be able to recognize variants of a known worm, but not a completely new worm. Even if antivirus vendors can analyze a new threat and make an update available in a few hours—no small feat—the update still needs to be distributed to many workstations and servers. It is also prudent to perform at least a rudimentary test of each new update before deploying it to ensure that it does not cause system crashes or other issues. Antivirus signature updates are often released multiple times a day, making it even more challenging to keep systems current. Even in a best-case scenario, there is still a sizable window of opportunity for a new malicious code threat to successfully infect systems before antivirus software can stop it. Accordingly, organizations need to use multiple security controls to stop new malicious code threats.

Several years ago, worms such as *Love Letter* spread through E-mails that used a fixed subject, attachment name, and message body. This allowed organizations to block infected E-mails in a matter of minutes by filtering on these simple characteristics. It was also easy for users to identify infected E-mails and avoid executing their attachments. Many recent worms have been much harder to detect and stop. Today's mass mailing worms typically use dozens or hundreds of E-mail subjects, file attachment names, and message bodies, or even generate them at random, so it is often impractical or impossible for E-mail servers to filter them. Some network service worms are difficult to detect once they infect a host because they only exist in memory; antivirus products that look for infected files fail to identify them. This underscores the need to have multiple security controls in place to stop malicious code, because no single control is capable of detecting and stopping every malicious code threat.

A layered approach to preventing infections

An effective defense against malicious code requires the combined efforts of several types of security controls. As discussed in the following sections, the necessary controls can be divided into three layers: network, host, and user.

Network layer

The role of the network layer is to prevent malicious code from reaching hosts by examining and blocking suspicious network-based activity. Security controls at the network layer typically include firewalls, routers, and other packet filtering devices; antivirus servers; and content monitoring and filtering services (e.g., E-mail messages, Web activity). The most important practices involving the network layer are as follows—

- **Use antivirus servers to identify and block malicious code**—As discussed earlier, antivirus software can be very effective at stopping known malicious code threats, particularly mass mailing worms. (Network-based intrusion prevention software may also be effective at stopping certain types of malicious code.) NIST SP 800–61 recommends deploying antivirus software at the network layer (in addition to the host layer) and keeping the antivirus software as current as possible. DISA's Network Infrastructure STIG requires antivirus checks to be performed on all incoming packets by firewalls or associated antivirus servers.
- **Only permit necessary network activity**—Configuring firewalls and other packet filters to restrict which protocols and ports may be used can be effective in preventing many network service worms from entering an organization's network, as well as stopping infected hosts on the network from spreading worms to other hosts. NIST SP 800–61 recommends this practice for malicious code prevention and containment. DISA's Network Infrastructure STIG specifies many protocols and ports that should be blocked at the network perimeter, and also states that packet filters

The best practices presented in the rest of this article are based upon the following sources—

- NIST Special Publication (SP) 800-61, Computer Security Incident Handling Guide. This contains recommendations for preventing all types of incidents, including specific guidance on malicious code infections. The best practices in SP 800-61 reflect the experiences of commercial, educational, and governmental organizations. The recommendations in NIST SP 800-61 tend to be somewhat general because they are intended to be applicable to a wide variety of environments.
- DISA's Security Technical Implementation Guides (STIG). The STIGs provide specific minimum requirements for network, host, and application security in DoD environments. Although the STIGs are not specifically directed at malicious code, following the requirements in the STIGs can be very effective at preventing many malicious code infections. Relevant STIGs include the Network Infrastructure STIG, the Desktop Application STIG, and the Windows STIGs.

such as routers should be configured to deny all network activity that is not expressly permitted.

- **Block suspicious files based on file extensions**—A simple but effective technique is to configure E-mail servers or filters to block E-mail file attachments that are likely to contain malicious code, such as files with .exe file extensions. This prevents many unknown mass mailing worms from reaching hosts, but could inadvertently block legitimate activity. NIST SP 800–61 recommends that organizations identify and block file attachment types that are potential carriers of malicious code if they are not necessary. DISA's Network Infrastructure STIG reflects this practice, listing several dozen file extensions to be blocked or otherwise filtered by E-mail servers.

Host layer

The role of the host layer is to reduce the likelihood that malicious code that reaches a host can infect it. This is done through two efforts: eliminating vulnerabilities or weaknesses that could be exploited, and examining and blocking suspicious activity. Host layer controls affect operating systems and applications, such as E-mail clients and Web browsers. As described below, good practices at the host layer include patching systems and applications, hardening hosts, using antivirus software, limiting the use of file sharing applications, and configuring applications more securely.

- **Keep systems and applications patched**—This eliminates known vulnerabilities that could be exploited by malicious code. It is becoming increasingly challenging to deploy patches quickly enough to prevent infections—the typical time from public announcement of a vulnerability to the release of malicious code has dropped from several months to weeks. The *Witty* worm was released less than three days after the vulnerability announcement. (Even if a patch is not deployed before malicious code hits, it will be needed afterward to prevent cleaned systems from becoming reinfected.) NIST SP 800–61 recommends, and the DISA STIGs require, that operating systems, Web browsers, E-mail clients, and other applications be updated with new patches in as timely a manner as possible.

- **Keep hosts hardened**—Besides keeping hosts fully patched, administrators should harden hosts' operating systems to limit the possible methods of attack and reduce the impact of successful attacks. For example, the DISA STIGs require that certain unneeded services be disabled or uninstalled, and separate user accounts with limited rights be used for regular system use. Disabling unused services prevents malicious code from exploiting current and future vulnerabilities in those services. Limiting user privileges can thwart malicious code that requires administrative-level rights to exploit a vulnerability. In addition to these practices, NIST SP 800–61 also recommends eliminating unsecured shares on systems, using personal firewalls on all systems, and configuring personal firewalls to block unauthorized incoming connections. All of these measures are effective at stopping many network service worms.

- **Use antivirus software to identify and block malicious code**—Each host within an organization should use antivirus software to identify attacks against the operating system and commonly targeted applications, including E-mail clients and Web browsers. Having antivirus software at the network layer is insufficient; for example, a virus may enter a system through removable media or through a local network segment not monitored by network layer antivirus services. NIST SP 800–61 recommends using antivirus software for workstation and server operating systems, as well as application server and client software when possible. The DISA STIGs require antivirus software to be installed and maintained on servers, workstations, and PDAs, and to update antivirus signatures at least weekly, preferably daily. (In addition to antivirus software, host-based intrusion prevention software may also be effective at blocking certain types of malicious code.)

- **Avoid the use of software with personal file transfer capabilities**—Examples of such software include public instant messaging and peer-to-peer music sharing services. Such services are often used to transfer files containing malicious code, typically disguised as benign content. The DISA Desktop Applications STIG requires that public instant messaging services and peer-to-peer file sharing programs not be used; NIST SP 800–61 also recommends limiting their use. The use of such software can be identified by checking systems, monitoring network activity, and blocking the use of certain protocols and ports.

- **Configure E-mail clients and Web browsers to behave more securely**—For example, NIST SP 800–61 suggests that E-mail clients could be configured not to open or run file attachments automatically, and Web browsers configured to limit mobile code execution. DISA's Desktop Application STIG has similar requirements, such as prompting the user to confirm opening certain file attachments and configuring systems to open some types of files (such as mobile code scripts) with a text editor instead of executing them.

User layer

The role of the user layer is to reduce the likelihood that users will trigger infections on the hosts they use. The main component of the user layer is educating users as to the requirements of the organization and the best practices they should follow in terms of system and application use and security. For example, a policy that forbids the use of public instant messaging services may be more effective if users are made aware of it and reminded periodically.

Providing training regarding safe E-mail attachment handling practices may be effective in reducing (but not eliminating) the infection rate for new mass mailing worms. If prevention mechanisms are unsuccessful in stopping malicious code from reaching users, it is very likely that an incident will occur because of the sheer numbers of people and systems involved. Imagine that thousands of

continued on page 30…

The Future of Network Intrusion Detection

by Abraham T. Usher, CISSP

Most commercial and government organizations have firewalls and anti-virus products to protect their infrastructure. Why are intrusion detection products necessary? There are two primary reasons —firewalls are not a complete solution, and not all malicious network traffic originates outside of the firewall.

Network firewalls in isolation are not adequate for maintaining security within Internet Protocol (IP) networks. All firewalls are subject to being bypassed by traffic that has been forged with false packet header data; additionally it is possible to "tunnel" malicious traffic through existing protocols. Tools such as httptunnel can encapsulate bi-directional Transmission Control Protocol (TCP) network connections and hide them in HTTP requests, effectively hiding communications in a covert channel. [1]

Figure 1 depicts a firewall blocking incoming telnet traffic that is attempting to connect to a telnet server using port 23. In this case, the traffic is intercepted by the firewall and blocked because port 23 is not allowed for use by inbound connections.

Figure 2 depicts a similar scenario, only this time the telnet session is piped to an *httptunnel* client that transmits the session via HTTP commands through port 80, to an *httptunnel* server that decodes the commands and forwards them to the telnet server.

There is no easy way to use firewalls to prevent attacks that are tunneled through other protocols. With few exceptions, most firewalls examine IP connections and valid IP state—they do not perform monitoring of application level payloads. Intrusion detection systems provide an important additional layer of defense against tunneled attacks.

Another shortfall of firewalls is that although they can stop incoming and outgoing traffic based on their policies, they cannot enforce rules on traffic between hosts within an internal local network. In order to have visibility of events occurring between hosts on the same internal network, intrusion detection systems are required.

There are two general categories of Intrusion Detection System (IDS) technologies—those that monitor individual hosts, and those that monitor networks. Additionally, there is a relatively new push towards intrusion prevention system (IPS) technologies that go beyond merely detecting events and attempt to prevent intrusions in near-real time. This article focuses on the evolution of network IDS into network IPS.

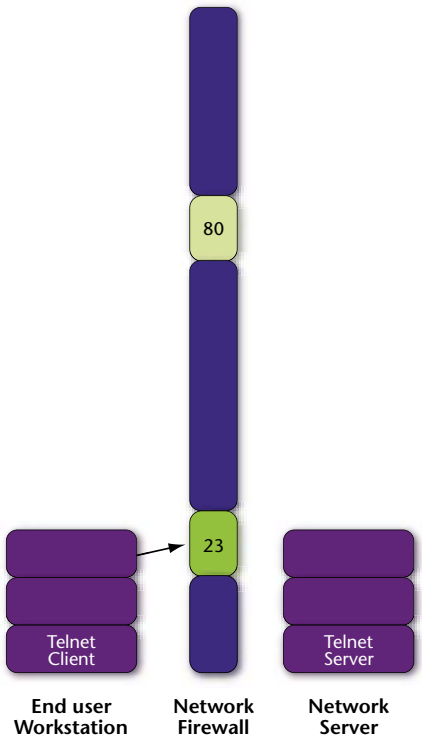

Figure 1: Firewall blocking telnet traffic

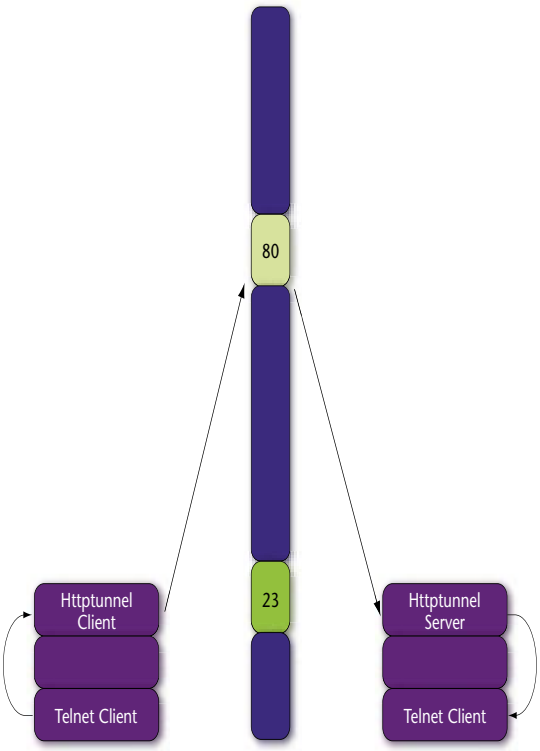

Figure 2: Firewall permitting traffic tunneled through port 80

The need for intrusion prevention

The OODA loop (Observe, Orient, Decide, Act) is a famous framework developed by John Boyd to depict the four critical actions that determine how rational actors respond to events. [2] The OODA loop is a useful construct for examining the importance of intrusion prevention capabilities (see Figure 3, page 12).

Consider an organization with a standard IDS. If an attacker performs a network scan and finds several vulnerable systems, he may launch an exploit a few seconds later. A good IDS would alert the system administrator through short message service (SMS) or E-mail within one second of the exploitation taking place (observation). Even a very experienced administrator would need at least five seconds to consider the incoming alert and its implications (orient). The administrator would need several seconds (at least two) to consider what to do next (decide). Finally, actually responding to the intrusion activity by implementing a change to the network configuration would take several seconds, at least five (act). In this very optimistic scenario, a seasoned system administrator would complete all stages of the OODA loop in thirteen seconds (see Figure 4, page 12). Although this sounds pretty fast, in the world of computer systems and automated attack scripts thirteen seconds is an eternity.

To reduce the defenders OODA loop, IPS introduces defensive capabilities that occur in computer time as opposed to human time. The entire cycle of observing, orienting, deciding, and acting on incoming network traffic is processed by the IPS (see Figure 5, page 12). In this example, each phase of the OODA loop requires 0.1 seconds of time by the computer for a total of 0.4 seconds. This is an improvement of over 3,000% when compared to the best-case human response time of 13 seconds.

In reality, the processing time could be substantially less than 0.4 seconds (providing a near real-time response capability). Of course when automated IPS technologies are implemented the critical “decide” phase is taken away from human actors and implemented by the policy of the IPS. This can be either an advantage or a disadvantage, depending on the system context and complexity of decisions being made.

Methods for achieving intrusion prevention

There are several basic mechanisms for achieving intrusion prevention capabilities on IP networks once unauthorized activity is detected. The main five methods include—packet spoofing, shunning, rate limiting, traffic reflection, and inline traffic analysis.

Packet spoofing occurs when the IDS forges a TCP RST (reset) packet to tear down the offending network connection. This method may not be successful in all cases, as

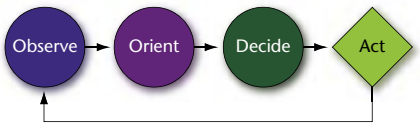

Figure 3: The OODA loop

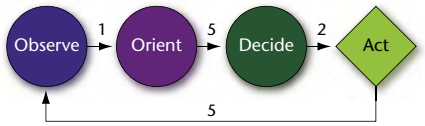

Total time to complete OODA loop = 13 seconds

Figure 4: Best-case human response time to an incident

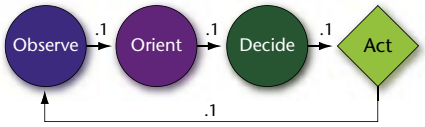

Total time to complete OODA loop = 0.4 seconds

Figure 5: Computer response time faster than human response time

some exploits only require a single packet, and some transactions are so fast they can occur before the connection can be broken.

Shunning is a method where an IDS dynamically reconfigures perimeter security (e.g., router access control list or firewall policy) in response to incoming malicious traffic. This method is also limited in its efficacy as reconfiguring network devices may take several seconds, which is ample time for many automated exploits to achieve their objectives.

Rate limiting is a special form of shunning where anomalous connections are downgraded to receive less available bandwidth. This technique does not prevent intrusions, but it may deter attackers enough to seek more attractive targets.

Traffic reflection is a process of blocking attack traffic and sending an exact copy of the attack back to the sender. This method of intrusion prevention is on shaky ground legally and is agains DoD policy, but is available in certain commercial products. [3]

Inline traffic analysis refers to a configuration where the IDS is deployed inline, similar to a firewall. The IDS receives traffic and performs an evaluation. Good traffic is forwarded to its destination and bad traffic is recorded then dropped. This class of IPS implementation appears to be the most promising of any method for use in production network environments. However, this method can be quite risky as an attacker may use it to invoke a denial of service against an authorized user by spoofing attack traffic with the authorized user's IP address. Care must be taken when deploying IPS to ensure legitimate users and software agents will not be blocked. Also, if an IPS ceased operation it should be configured to fail open (if availability is critical) or fail closed (if confidentiality is critical).

Drawbacks of IPS

IPS products show some promise through their abilities to intercept and stop malicious traffic. However, these products introduce many risks—

- IPS products are not as scalable as IDS products
- If an IPS fails, it could create a single point of failure for an entire network segment
- IPS may cause self-inflicted denial of service (DoS) and block legitimate user traffic

Although IPS products show promising results for low bandwidth (100 Mbps) networks, they may not be mature enough for deployment in high bandwidth, high availability networks. Also their "prevention capabilities" present an Achilles heel whereby legitimate users and agents could be blocked from accessing a network due to spoofed malicious traffic or system misconfiguration.

Due to the inherent limitations of IDS, Gartner Group released a report in June 2003 calling IDS a "market failure." Gartner subsequently recommended IPS and deep inspection firewalls as more promising technologies. IDS technology does have some considerable drawbacks and complexity. However, the use of IDS technology provides critical insight into identifying security violations on computer networks and is crucial for the effective implementation and compliance monitoring of security policies.

The shrinking window from vulnerability to exploitation

The window of time between a security vulnerability being publicly announced and subsequently being exploited by malicious code is shrinking at an alarming rate (Figure 8). Most information technology departments take more than 30 days to test and apply newly released software patches. Recent malicious code outbreaks demonstrate that the old paradigm of "react and patch" is ineffective. For example, the time period from the announcement of Microsoft Security Bulletin 04–011 "Local Security Authority Subsystem Service (LSASS) Vulnerability" to its exploitation by the *Sasser* worm was only 18 days. As they become more mature, intrusion prevention systems may provide a valuable security mechanism for countering the threat posed by the shrinking window of vulnerability discovery to vulnerability exploitation.

Future trends

Three general trends seem likely as organizations seek to secure their enterprise networks through new technologies—

- Integration of security components
- Active response capabilities
- Return to protection of network end points

IDS products are becoming increasingly integrated with other security systems. Rather than being stand-alone security tools, IDS products will exchange data with other systems (firewalls, vulnerability assessment tools, network management consoles). Historically, incident response staff members have spent a large amount of their time trying to figure out which IDS alerts are genuine and which are false alarms. By integrating IDS sensor data with known system

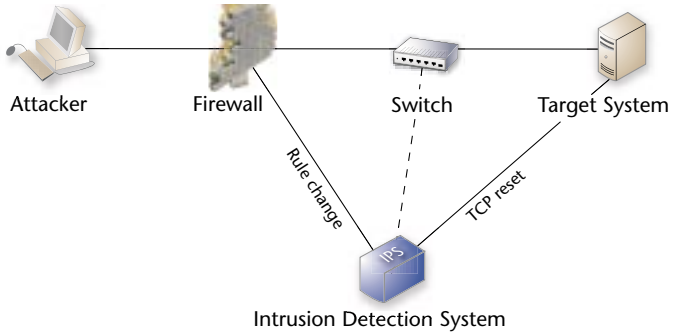

Figure 6: Traditional IDS with TCP reset capability

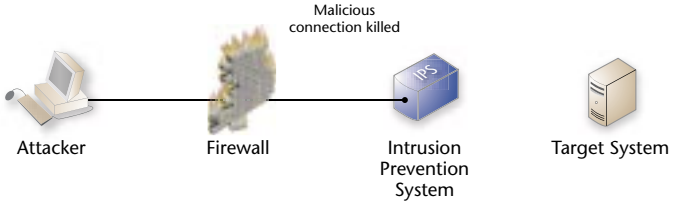

Figure 7: Intrusion prevention system blocking traffic in-line

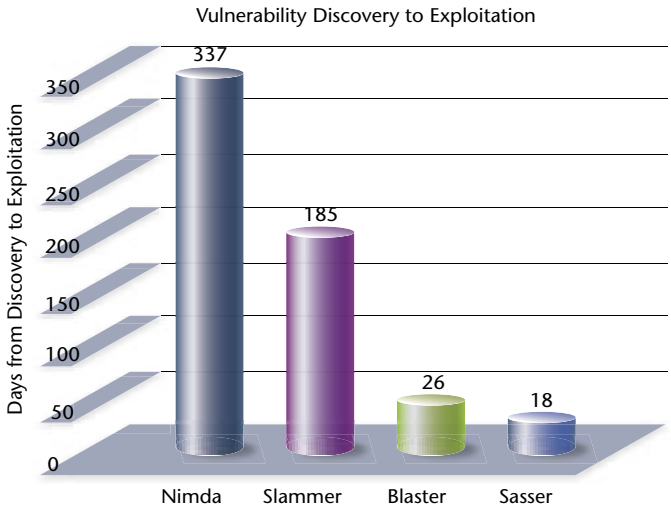

Figure 8: The narrowing window of time from vulnerability to exploitation

vulnerability data, the number of false positives can be substantially reduced so that security staff can focus on the most important events.

Many IDS are now including active response capabilities to compliment the passive response measures that classical IDS have. These active response capabilities include TCP resets (for closing hostile sessions), router and firewall update messages (to reconfigure firewalls against IP addresses that are sending malicious traffic), and in-line packet evaluation (forwarding acceptable traffic, and dropping unacceptable traffic). In some cases, products even include "hack back" capabilities against malicious traffic—this practice is legally questionable. [4] With any system that has active response capabilities, care must be taken that legitimate users are not prevented from using computing resources that they require.

It is easy to make IP networks work—it is difficult to make them work well. Coupled with the increasing complexity that modern information security adds to existing network architectures, many organizations are considering focusing their information protection efforts on network end-points such as user workstations, application servers, and network devices. As network perimeters are increasingly exposed to additional threats through virtual private networks and wireless access points, increasing the security measures for high-value hosts will become a necessity. ■

About the Author

Abraham T. Usher, CISSP

Abraham T. Usher graduated from the U.S. Military Academy in 1996 with a B.S. double major in Modern Standard Arabic and German language studies and he also received a M.S. in Information Systems from George Mason University. Mr. Usher is a Certified Informational System Security Professional (CISSP). and may be reached at abe.usher@sharp-ideas.net.

References

1. HTTP Tunnel. http://www.nocrew.org/software/httptunnel.html
2. Boyd, John. "Boyd and Military Strategy." http://www.d-n-i.net/second_level/boyd_military.htm
3. Symbiot. "Symbiot Intelligent Security Infrastructure Management Systems." http://www.symbiot.com/
4. Gaudin, Sharon. "Plan to Fight Back Against Hackers Causes Stir." Security Planet. March 2004. http://www.internetnews.com/article.php/3327391

IPv6
The Next Generation Internet Protocol

by Matthew Warnock

The Internet was born as a project to connect Department of Defense (DoD) computers together over long distances. This network, created in the late 1960s, was called the Advanced Research Projects Agency network (ARPAnet) and it connected several research centers together. As the need for connectivity between computers on this network grew, so did the need to connect ARPAnet and other networks across the globe. It was necessary for a standard to be created to connect any kind of computer (IBM, Unisys, etc.) over any kind of medium (cable, radio, satellite, etc.) since there were many computer architectures and network connections available at the time. In the 1970s, a new protocol called the Internet Protocol (IP) was proposed that could do just this. The IP, the same IP in transmission control protocol/internet protocol (TCP/IP), was adopted a few years later.

The version adopted in the early 1980s was called Internet Protocol version 4 (IPv4). At that time, no one had any idea that an Internet boom would be occurring not much more than a decade later. In the mid-1990s every business, organization, and individual was getting "online." This growth occurred at an average of 10x every two to three years since the beginning of the 1980s. The structure of the Internet has not changed very much, and that can attest to the resilience of IPv4.

Close to 25 years and 2.5 billion online users later, the Internet is demanding an upgrade. The Internet Engineering Task Force (IETF) laid out the structure for this update in 1994 in a standard called Internet Protocol version 6 (IPv6) or IPng for next generation. It was adopted in 1998 as the standard for the next generation of the Internet Protocol, but why do we really need it?

Inadequacies of IPv4

We have a little more than 2 billion usable Internet IP addresses available, 6 billion people in the world, and many more computers than people. Something has to give. Towards the end of the 1980s, it was realized that we are running out of these IP addresses. IPv4 uses 32-bit addressing for its connections. Each octet of an IP address is 8-bits, (i.e., 208.254.0.16).

Four of these 8-bit octets add up to a 32-bit address. Theoretically, there should be 2^{32} addresses available, or 4,294,967,296, but the addresses were distributed inefficiently. About 270,000,000 were allocated for internal network usage like private, loop-back, and multicast IP addresses. About 1,500,000,000 addresses are "reserved" by the Internet Assigned Numbers Authority, IANA, leaving about 2,500,000,000 available for use—most are already being used.

One quick-fix solution to the problem of limited IP addresses is Network Address Translation (NAT), which allows several computers on a network to share a public IP address. Each user on the private network gets a private IP address and they share one public IP address. While this solution works well to a degree, it creates problems for many applications that require direct access to the Internet. NAT only allows for outbound traffic and will not allow for inbound connections to the private network. Services on the private network, like file transfer protocol servers, are usually unavailable. An example of this is shown in Figures 1 and 2. In Figure 1 (see page 15), each workstation and server has its own public IP address and connections can be made on any network to any server. A problem arises when a private network is established, shown in Figure 2 (see page 16), when each computer does not have a unique public IP address, but instead has only one address. Devices within the private network can reach the server—however, any devices outside of the network cannot reach the server. Since NAT does not allow incoming connections, security increases slightly because it is harder to penetrate the private network.

It may be more difficult to enter the private network using NAT, but overall IPv4 is not very secure. Internet protocol security (IPSec) compatibility is not mandatory and there is no standard for encryption. IPv4 source addresses can be easily spoofed and connections are not always traceable.

Benefits of IPv6

Because of IPv4's resilience, it has served its purpose well and for that reason, the next generation Internet

Protocol is largely based on IPv4. IPv6 uses the same kind of header, just tweaked a bit, and several features have been added while obsolete ones have been removed.

Increased address space

The most notorious feature of IPv6 is the increased address space. This will probably be what fuels conversion from IPv4 to IPv6 for most of the world. The new 128-bit addressing provides many more available IP addresses. The number of usable IP addresses has been estimated to be around 320 trillion. While it would be nice to say we will never run out of IP addresses, the same thing was said in the early 80s when the colossal 32-bit addressing was chosen over the more reasonable 16-bit. Besides having a plethora of IP addresses available, a plan for more efficient allocation of the addresses is in place.

On an IPv6 network, every computer can have its own IP address, eliminating the need for NAT and opening the door for newer technologies. Using the IPv4-based Internet, only computers are considered available to the Internet. The advent of IPv6 will allow for mobile devices, game consoles, organizers, and other non-computer devices to have their own IP address.

Security

DoD's interest in IPv6 is not limited to increased address space but in the other features like security. IPv4 makes IPSec optional, but IPv6 makes its compliance with IPSec mandatory. An IPv6 host must support IPSec's Authentication Header (AH) and Encapsulated Security Payload (ESP) protocols. AH and ESP provide authentication and data integrity and ensures the destination host knows the data origin. This is a major benefit to DoD because currently IPSec is not mandatory in IPv6.

Configuration options

Some features that will improve network administration are the configuration options. The Autoconfigure feature allows a computer to be setup on a network without the use of a Dynamic Host Configuration Protocol (DHCP) server. The client uses a combination of the subnet prefix, which

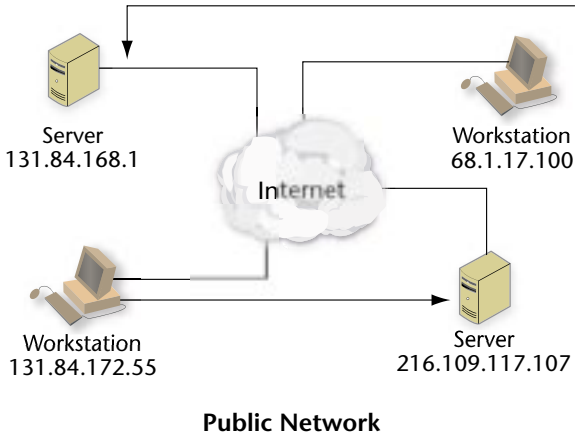

Figure 1: Direct connections

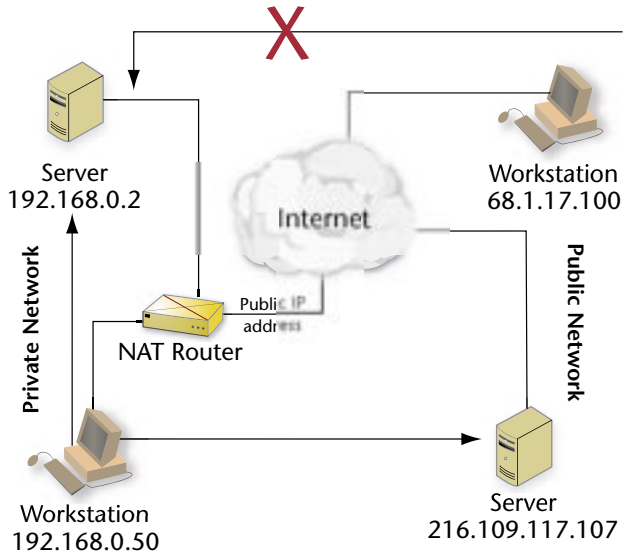

Figure 2: Network address translation

is a network identifier that all clients on the network share, and the local identifier, specific to the client machine (see Figure 3, page 16). Often, the local identifier is a variation of the 64-bit Media Access Control (MAC) address. Once the 128-bit address has been calculated, the client will send messages to the router to make sure the address has not already been delegated. This is called Duplicate Address Detection (DAD). The client now has a unique IPv6 IP address because of the Autoconfigure feature.

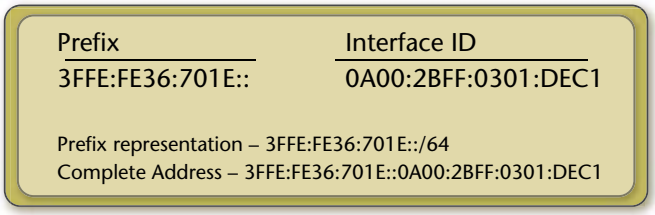

Figure 3: Structure of an IPv6 address

Routing

IPv6 allows for more efficient routing as well by eliminating the large tables used by IPv4 routers. Instead, a hierarchy was created to allow only the routes of the connections to each router to be known. This hierarchy is based on geographic connections so the data path is efficiently sent around the world, increasing the speed of communication.

Simpler header structure

While the length of the addresses in IPv6 is four times larger than the addresses in IPv4, the IPv6 header, which contains the source and destination addresses, is only twice as large. Much of the unnecessary information was removed from the IPv4 header to streamline it and optimize it for limited overhead. The router only needs to process essential information from the header, and this ensures that IPv6 packets are efficient enough to travel on slower networks as well as fast ones.

Differences between IPv4 and IPv6

IPv4 is a very impressive protocol and that is why it has lasted for over twenty years and why IPv6 is largely based on IPv4. You will not see some of the differences between IPv4 and IPv6 unless you look under the hood of each protocol. See Figures 4 and 5 on page 17 and examine the differences between IPv4 and IPv6 header structure. IP protocols are layer 3, or network layer, protocols, which is the mechanism that gets data packets from source to destination address. IPv4 and IPv6 can work side by side on the same network, but are not compatible.

IPv4 header contents—from RFC 791

```
Version – 4-bit – Version of the protocol
(version 4)
Internet Header Length – 4-bits – length in 32
bit words
DS byte – 8-bits – Priority delivery value
Type of Service – 8-bits – sets precedence to
different packets
Total length – 16-bits – measured in octets
Identification – 16-bits – created by sender to
identify fragments of data
Flags – 3-bits – information about fragmentation
Fragment offset – 13 bits – measured in
fragments, where this fragment belongs in the
datagram
Time to live, TTL – 8-bits – decrements until the
it arrives at its destination, or is terminated
Protocol – 8-bits – indicates next level protocol
used in the data portion
Header checksum – 16-bit – error checking for the
header only
Source address – 32-bits - Address of sender
Destination address – 32-bits - Address of
intended recipient
Followed by data
```

IPv6 header contents—from RFC 2460

```
Version – 4-bits – Version of the protocol
(version 6)
Traffic class, DS byte – 8-bits – Priority
delivery value
Flow label – 20-bits – Router handling for
sequence of packets
Payload length – 16-bits – length of data
Next header – 8-bits – Type of the next header
Hop limit – 8-bits – Number of "hops" before
packet is terminated, similar to TTL in IPv4
Source Address – 128-bit – Address of sender
Destination Address – 128-bit – Address of
intended recipient
Followed by data
```

The Internet Header Length (IHL), type of service, total length, identification, flags, fragment offset, protocol, and header checksum have been removed. The Time to Live (TTL) has been updated and replaced with hop limit. payload length has been added. While the IPv6 header is larger because of the address size, it has been optimized for efficiency and is only twice as large.

Message types

The IPv6 protocol utilizes three types of messages: unicast, multicast, and anycast (Figures 6, 7, and 8 on page 18). Only the anycast message is new. The unicast message is the basic one-to-one communication. The message is hierarchal because the message path is based on the global routing prefix, subnet ID, and interface ID. If a message is sent to another computer in the same subnet, it never has to leave that subnet. The same is true if a message is sent to a computer in the same network.

The second type of message is the multicast message or one-to-many communication. Just like the unicast messages, the multicast message can be sent as a site local message, which is on the same network. The multicast message can also be sent as a link local message, which is on the same subnet; however, multicast messages cannot be sent globally. A message will be sent to a subnet, and only certain devices will listen for the multicast messages. Because the multicast messages can be sent to a subnet of your subnet or network, these addresses can be reused on other sites and local networks.

The new type of message is the anycast message—rightly named because it can be sent to any one machine that was intended to receive it. This is like a multicast message because any machine on a list can receive it, but only the recipient with the most efficient route is chosen.

Network administrators

The conversion to IPv6 must start at the edges of the Internet and continue into the core. Because of this, system administrators will see most of the technical side of the rollout. The heart of IPv6 is in the connection hardware of a network, such as the routers, layer-3 switches, and firewalls. Manufactures have already started shipping

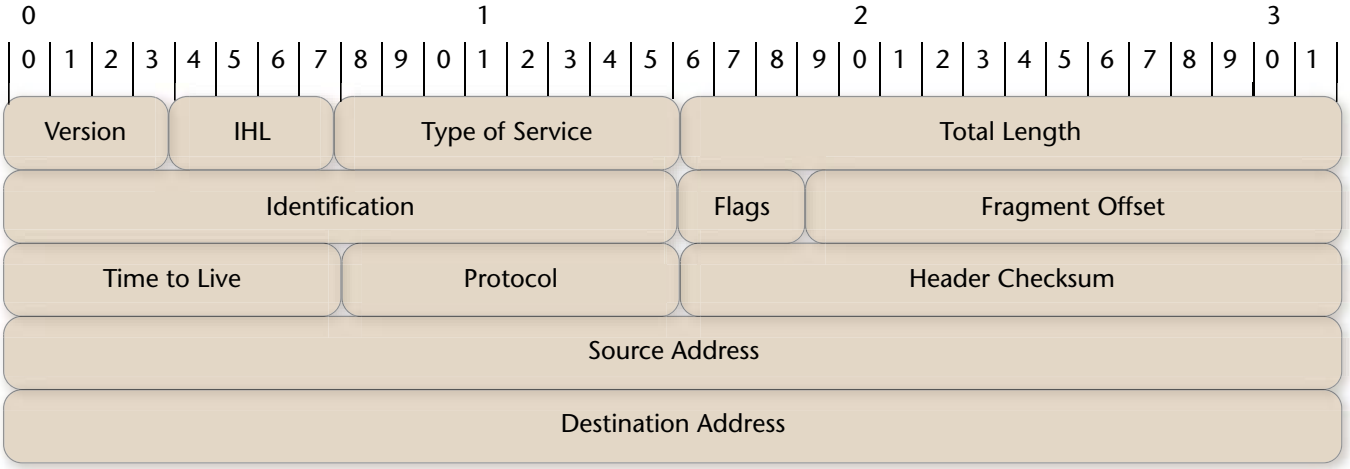

Figure 4: IPv4 header – From RFC 791

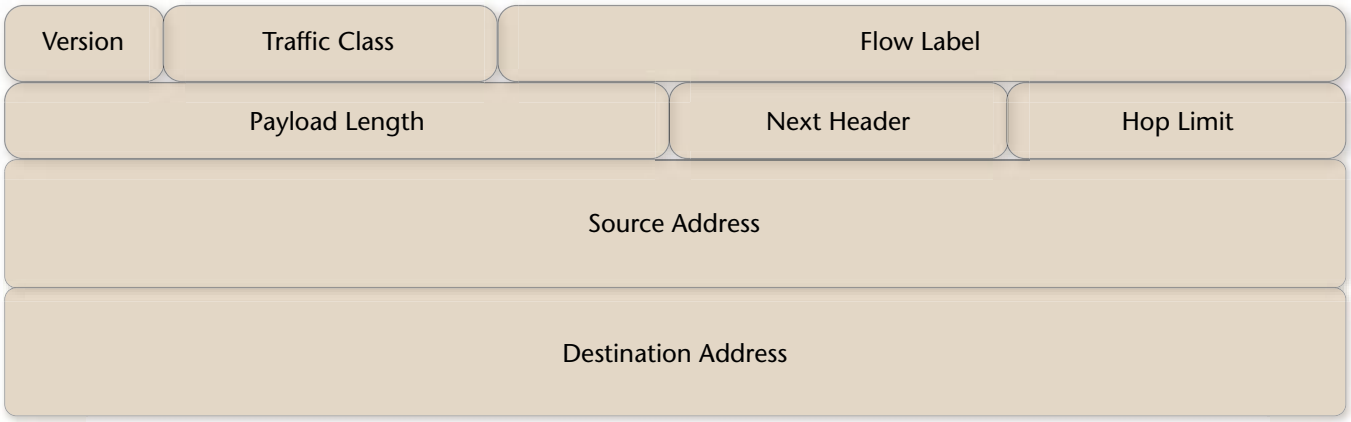

Figure 5: IPv6 header – From RFC 2460

hardware that is IPv6 capable, so the devices that you are using may already work with IPv6, or they can be upgraded to be compatible. Any hardware that does not look at the packet, such as switches, network cabling and network interface cards (NICs) does not need to be upgraded because it does not care what protocols pass through it.

Because upgrading to IPv6 requires changes to applications, software must also be updated to be IPv6 capable. This will include workstation and server Operating System (OS) upgrades to work with the new protocol, as well as upgrades to application and server software such as Exchange, Internet Security and Acceleration (ISA), and Domain Name Server (DNS). Currently, Windows 2003 Server near-fully supports IPv6 as well as Windows XP to a degree. Windows XP service pack 2 (SP2) promises more support for IPv6, but we will not see total compatibility in OSs until the new protocol is more fully implemented on networks. Exchange and ISA server software are slated to support IPv6 in their newest major releases. The Windows DNS server already fully supports IPv6.

Users

Users will experience the least amount of effort necessary to be IPv6 capable. Most companies, agencies or home users upgrade their OS every few years and most operating systems, including Windows 2000, XP, MacOS X, SunOS, Linux and others, support IPv6 in some fashion. Once the OS is supported, more applications will become available as well. Internet applications such as Internet Explorer will need upgraded to support IPv6. Internet Explorer 6 already supports IPv6 so it is just a matter of upgrading to version 6. Once a user's software is ready for IPv6, the user must be prepared as well. The common user does not run across IP addresses very often because Internet domain names have made it so easy to avoid them. If a user does want to use an IP address, they will see the address looks a little different. For starters, they are much longer. Users will also find that they do not have to be just numbers because they are represented in hexadecimal. Hexadecimal can represent a 16-bit number in four digits making the numbers visually smaller. 0–9 in decimal is 0–9 in hex, but 10–15 in decimal is A–F in hex. IPv6 addresses are written like this—

```
ABCD:ABCD:CCCC:DDDD:1234:5678:1111:2222
```

Every digit can be a number 0–9 or a letter A–F. There are eight sections of four hexadecimal digits separated by colons. These long addresses require a bit more effort to remember them, so to help with that, the standard, RFC 1924—A Compact Representation of IPv6 Addresses, is in place so that anytime there is a leading zero in an octet, the zero may be eliminated—

```
3ffe:2a00:0100:7031:0000:1
```

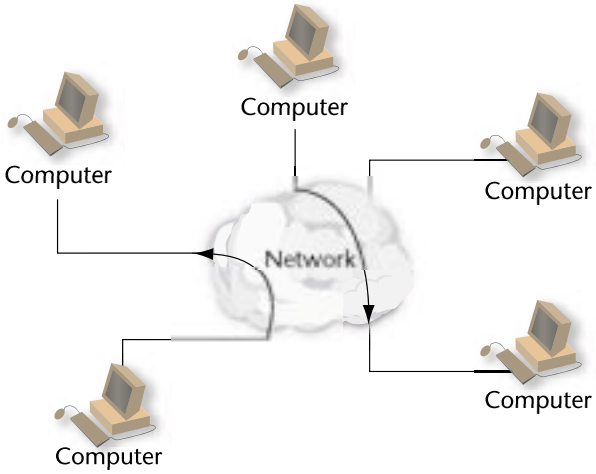

Figure 6: Unicast message

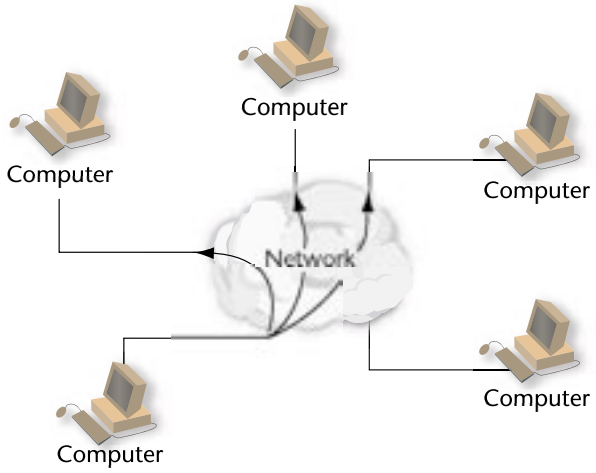

Figure 7: Multicast message

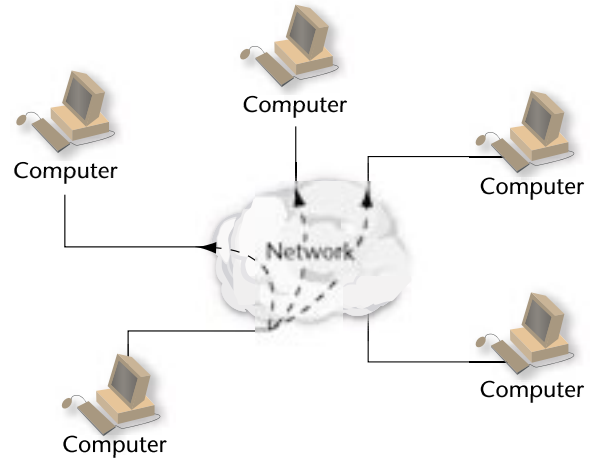

Figure 8: Anycast message

Becomes—

```
3ffe:2a00:100:7031::1
```

Users will find the uniform resource locator (URL) addresses containing an IP address looks a little different. To alleviate any confusion using colons and ports in URLs, the new standard of URL addresses uses brackets on the outside of the IP address—

```
http://[3ffe:2a00:100:7031::1]:80/index.html
```

Lastly, and probably most important, a home user must find a new ISP that supports IPv6. A few ISPs worldwide are IPv6 capable and this will increase as IPv6 develops.

Deployment

IPv6 will deploy worldwide over a relatively long period. Because a standard has been approved and networks have been set up, IPv6 is being deployed right now—however, it will require a lot of work to bring it together for the mainstream market. Currently, Asia is the furthest along in IPv6 development followed by Europe and then the Americas. As stated before, the transition must begin at the "edges" of the Internet and continue into the core, meaning that small network segments must convert to IPv6, and remain compatible with IPv4. DoD can implement IPv6 before the rest of the Internet and remain in service and connected to IPv4 (see Figures 9a, 9b, and 9c on pages 19–20).

To make this transition at a controlled pace, network connections must be upgraded to maintain backwards compatibility. IPv6 traffic must be able to travel over IPv4 networks, IPv4 over IPv6, or both at the same time. IPv6 traffic can tunnel through an IPv4 network by encapsulating the data to travel on the IPv4 network and decapsulating it when it gets to the IPv6 network (see Figure 10 on page 21.) These kinds of efforts require little cost or risk and are great methods for building temporary solutions but because of the extra encoding and decoding, should not be used for large amounts of data or used as a permanent replacement. Another way to let data travel over two different kinds of networks is to use translation where the headers are translated from one protocol to another (see Figure 11 on page 21.) IPv4 and IPv6 can coexist on the same network if dual-stack hardware is utilized, but this venture usually requires more effort and cost because of the addition of new hardware. Every type of transition method, dual stack hardware, tunneling, and translation, can be expected to be part of the rollout and methods to remain backwards compatible.

Any IPv6 solution set in place during the transition time must be very flexible. Some things can upgrade easily and some cannot, so make sure the solutions set in place can cover every activity on the network. The two IP protocols are very similar so this transition should be relatively easy. Reports from the U.S. IPv6 Summit in 2003 stated that DoD and universities received positive feedback about their IPv6 deployments. While hardware, software, and user techniques must be changed, the transition is not as difficult as telling a right-handed person to be left-handed

When can we expect to be on the IPv6 Internet? That is not easily answered. In Asia, you can purchase Small Office Home Office (SOHO) routers that are IPv6 capable

and Sony says every Internet-capable device they make starting in 2005 will be IPv6 ready. DoD wants to be IPv6 capable for the most part by 2008, but there will be pockets of IPv4 for a long time. IPv4 IP addresses are running out, so that will keep the IPv6 project moving.

DoD's deployment timeline

The Joint Technical Architecture (JTA) Development Group has set the standards for DoD's implementation of IPv6. These goals were laid out on June 9, 2003 in a DoD CIO memo—

- **September 30, 2003**—IPv6 address space allocated for DoD
- **October 1, 2003**—All network components purchased or acquired must be IPv6 capable
- **December 30, 2003**—IPv6 address space and naming convention described
- **FY 2005–07**—Segments of GIG convert to IPv6
- **FY 2008**—Near-full implementation of IPv6 on DoD networks

However, no current implementation of IPv6 in DoD other than separated test networks is allowed.

Besides becoming the standard protocol for the Internet, IPv6 will replace IPv4 on the Global Information Grid (GIG). This will make it necessary to upgrade the unclassified but sensitive IP Router Network (NIPRNET) the Secret Internet Protocol Router Network (SIPRNET) the Joint Worldwide Intelligence Communications System (JWICS) and any emerging DoD space and tactical communications that currently use IPv4 or will connect to the GIG or Internet.

Conclusion

The transition from IPv4 to IPv6 will not be automatic. It will not happen overnight, and it will not happen in the same manner worldwide. There is still much work to be done setting the standards around IPv6 and implementation will change the design of IPv6 as needed. Demand for this next generation Internet Protocol will fuel the transition; however, demand for legacy IPv4 products will ensure that we will see pockets of IPv4 networks throughout the world for a very long time. All forms of transition methods will be utilized to keep every activity world wide on task. Expect to see the beginnings of the IPv6 transition in the next couple of years in Asia and Europe and over the next five years in the Americas. ■

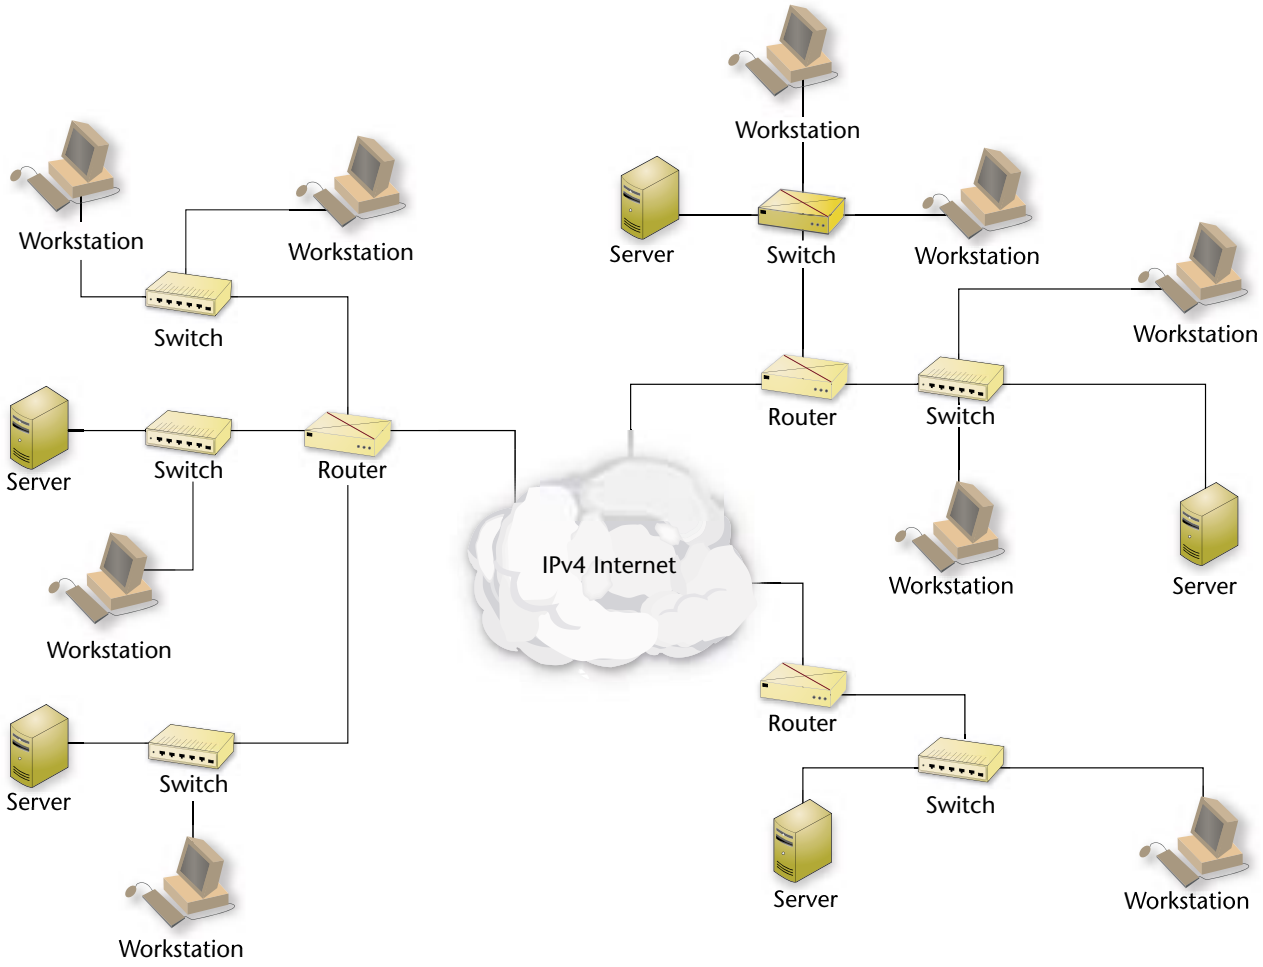

IPv4 Configuration

Figure 9a: IPv6 deployment stages

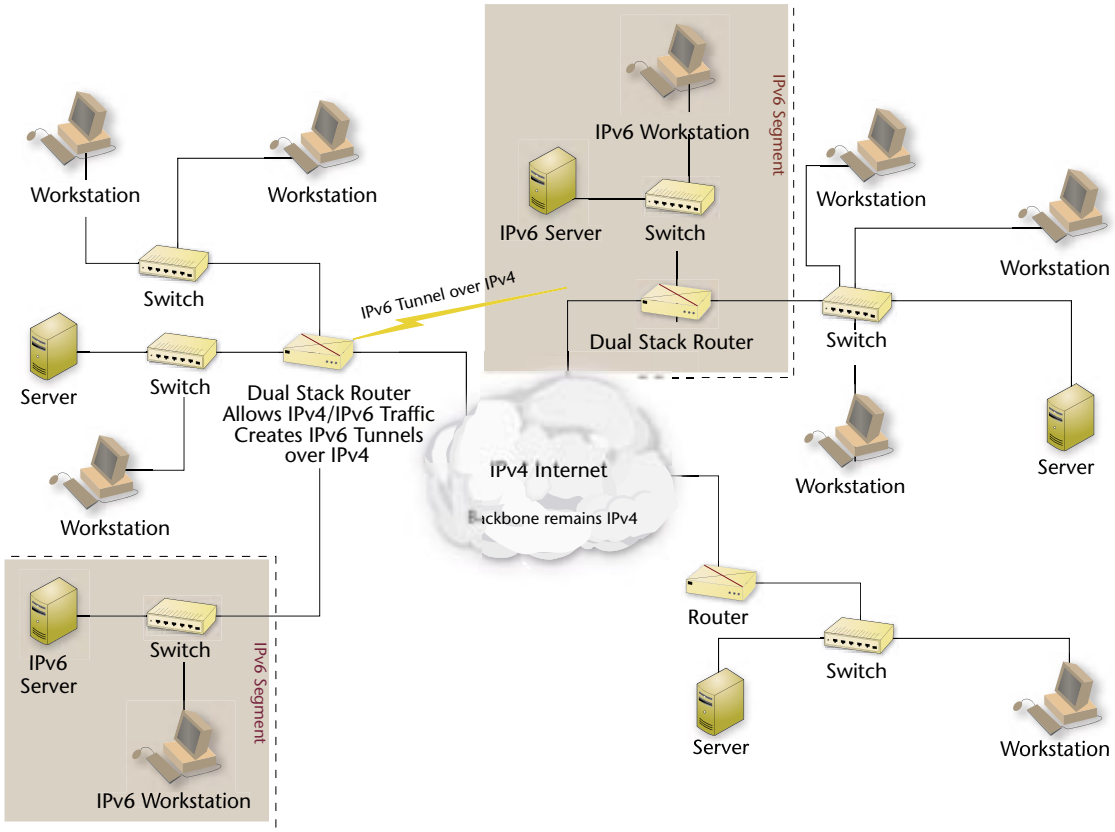

IPv4-based Networks with IPv6 Tunnels

Figure 9b: IPv6 deployment stages

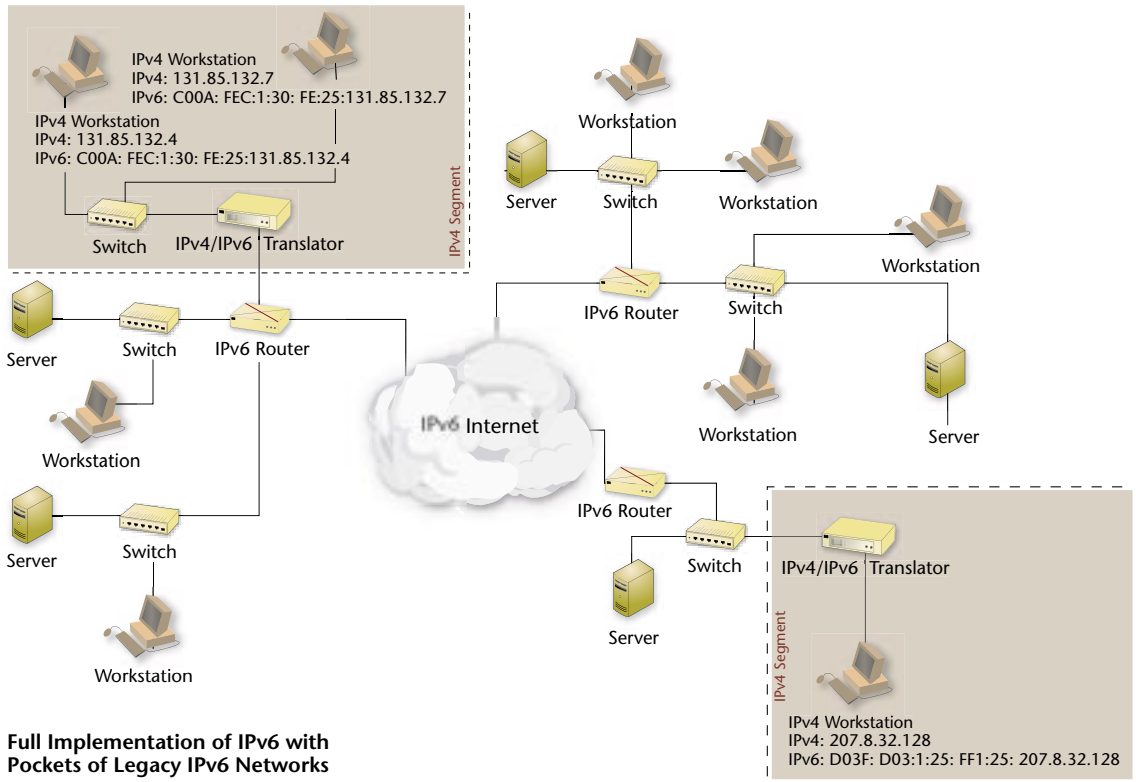

Full Implementation of IPv6 with Pockets of Legacy IPv6 Networks

Figure 9c: IPv6 deployment stages

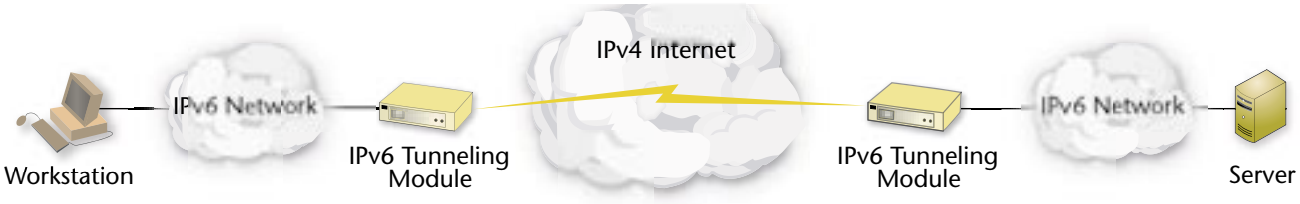

Figure 10: IPv6 Tunnel over IPv4

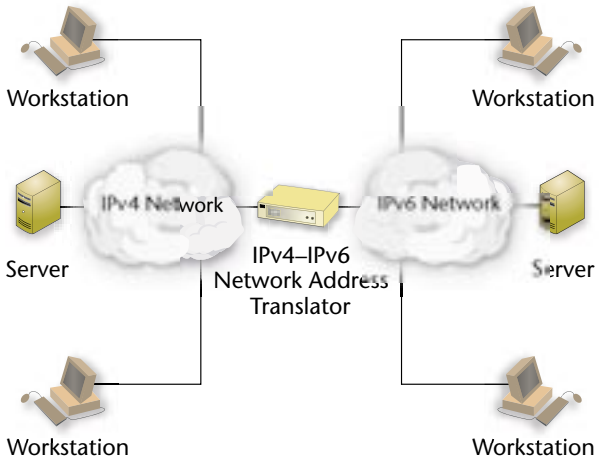

Figure 11: IPv4/IPv6 Translation

About the Author

Matthew Warnock

Matthew Warnock is an Information Assurance Specialist with the Information Assurance Technology Analysis Center (IATAC). He graduated from Pennsylvania State University with a B.S. in Electrical Engineering. His background includes assignments with the Defense Logistics Agency (DLA) in firewall and border protection support. Mr. Warnock may be reached at iatac@dtic.mi.

References

Hinden, Robert H., "IP Next Generation Overview" 14 May 1995.

"Internet Protocol Version 6 (IPv6)," DoD CIO Memo, 9 June 2003.

"Introduction to IP Version 6," Microsoft Corporation, March 2004.

"RFC 791: Internet Protocol, DARPA Internet Program Protocol Specification," September 1981.

"RFC 2460: Internet Protocol, Version 6 (IPv6) Specification," December 1998.

The Importance of High Quality Information Assurance (IA) Metrics

by Vivian Cocca, Steven Skolochenko, and Jonathan Smith

Information Assurance (IA) managers make risk-management and resource-management decisions in an information overload environment—they are inundated with information generated by the enterprise and must contend with numerous external data calls while attempting to keep up with the data associated with their own organizational responsibilities. This constant demand for IA management information carries with it the risk that data generators will become less precise when responding to managers' data calls, and the risk that managers will miss critical data when sorting through the volumes of data they receive. Awash in external reporting requirements, managers may sacrifice the collection of data they need to do their jobs well in favor of the collection of data they need to meet compliance reporting requirements. The Federal Information Security Management Act (FISMA) is a particularly notable (but certainly not unique) requirement that imposes an expensive and high profile collection on the Department of Defense (DoD) and its components.

FISMA reporting requirements are broad, and although they provide a general overview of the status of IA within an agency, they are not sufficient to effectively manage an enterprise-wide IA program. It is incumbent on the DoD to effectively and efficiently collect and report upward information required by statute and also to determine what metrics it needs to manage its IA program. That determination is not trivial—the DoD and its components must exercise foresight and good judgment in IA metric identification, collection, and interpretation. The faith that warfighters have in mission support information and information systems depends on effectiveness and efficiency with which the IA services are provided. Furthermore, the ability to resource and assess the success of the IA program is dependent on selection and use of "quality" IA metrics. The characteristics of a quality IA metric are the topic of the remainder of this discussion.

The role and value of IA metrics vary depending on one's organizational level. Later we will discuss the differences between, for example, the IA metrics needed by a system administrator and those needed by DoD's senior IA leadership. But first it is important to discuss what is meant by IA metrics. The term metric is often used to describe both the question asked and the data generated to answer it.

IA metrics are tools that support decision-making. Like experience, external mandates, and strategies, IA metrics are one element of a manager's toolkit for making and substantiating decisions. IA metrics are employed to answer three basic questions—

1. **Am I implementing the tasks for which I am responsible?** Consider the example of a program manager with responsibility for 250 IT systems. Among other things, that manager is responsible for the certification and accreditation (C&A) of those systems. A commonly used implementation metric for C&A is the percentage of systems accredited.

2. **How efficiently or effectively am I accomplishing those tasks?** Such IA metrics often answer more complex questions after an IA activity is fully implemented. For example, Federal law requires that C&A take place following a major system change. One might measure the efficiency of a C&A program by determining the time lag between each major system change and that system's renewed accreditation. Or one might measure the effectiveness of a C&A program by determining the number of accredited systems whose certification process included the creation of a system security plan.

3. **What impact are those tasks having on the greater DoD warfighting mission?** IA activities are initially selected with the belief that they will contribute to the DoD mission. After an activity is shown to be fully implemented and implemented well, managers must validate that the activity is delivering the expected benefit. These IA metrics are the most difficult to generate. A C&A process might be proved to have impact by showing that fewer interruptions or losses of data due to security incidents are experienced among correctly accredited systems than among incorrectly accredited or non-accredited systems.

Not all things measured or collected are IA metrics or components of an IA metric. These items are sometimes useful, but do not necessarily provide the insight or rigor expected of a quality IA metric. IA metrics are not—

- **Measures**—Data points such as "the number of system administrators" or "the number of systems" are measures. Useful and often necessary components of IA metrics, measures contribute to the meaningfulness of IA metrics. The major distinction is that a measure counts items that may not be the direct output of an IA activity.

- **Opinion polls**—Consider a survey of 250 system owners that asks, "How effective is the C&A process?" Responses may give a sense of the mood of staff, but will not provide insights into specific problems with a C&A process; nor will they substantiate possible corrective actions.

- **Project management milestones**—These are usually one-time events such as: "complete rollout of Internet Explorer Service Pack 2 to all Microsoft PCs." Once this task is done, it can be forgotten. Compare this to annual IA awareness training. While there is a discrete beginning and end to each year's training cycle, IA awareness training is a continuing requirement. Performance can be tracked over time to identify process improvements.

- **Performance Targets**—Consider the statement: "issue 90 percent of IA vulnerability announcements (IAVA) within one business day." This is not an IA metric. Rather, it describes a target of performance for the "time to IAVA issuance" metric. Performance targets should be excluded because targets may change from year to year.

As mentioned earlier, many IA metrics are collected because of an external data call. These collections usually address compliance with a particular law, regulation, or policy and may or may not be viewed locally as contributing to effective management of an IA program. However, since these must be collected, it is reasonable to consider them as sources of data for use in assessing progress against goals and objectives. Where a question is asked in a "low quality" way, an organization should answer as required—but it is often worthwhile and no more expensive to collect higher quality data or ask additional relevant questions.

Appending an organization's own IA metrics to a compliance-related questionnaire is a useful technique to take advantage of an activity that otherwise brings limited direct utility. But most useful metrics will be created directly by those responsible for an IA activity. These are often developed from scratch, but an under-utilized technique is to develop "compound" or "derived" metrics. Derived IA metrics take advantage of existing data points and process them to create new information. For example, a service might correlate information on IA security awareness training with information about the times and locations of successful system intrusions. The result might provide new insight into how awareness training does or does not strengthen the Department's ability to detect, resist, and recover from an attack.

There are eight characteristics of quality IA metrics—

1. **Strategic**—The DoD IA program has issued an IA strategic plan (DoD IA Strategy) that describes the overall mission, vision, goals, and objectives for the DoD IA program. All IA metrics should ultimately map to one or more of the five major goals of that plan. Consider the manager of a firewall program. Firewalls are not deployed for their own sake. They are generally used to resist attacks from external systems. That resistance is, in turn, part of an integrated DoD effort to defend systems and networks (strategic goal two). Computer net-

continued on page 32...

DEFCON 12 Security Conference Thoughts, Theories, and Comments

by Louis Lerman

In stark contrast to its history, the term "hacker" has developed a bad name. A dictionary definition of a hacker might read, "one who is proficient at using or programming a computer; a computer buff." However, a person on the street might relate a hacker to a criminal. It's important, therefore, to understand that "hacker" should not be used in the same sentence with "criminal," not at least because the author considers himself somewhat of a hacker. In fact, two weeks ago, he attended a conference in Las Vegas with many of his fellow hackers—some of whom are the brightest minds in the field of computer security.

DEFCON 12 took place July 30—August 2 in the Alexis Park Hotel in Las Vegas. For three days, lectures on many information assurance (IA) topics were delivered, with topics ranging from the technical, such as Paul Wouters' talk on Windows WaveSEC Deployment, to the political, with talks such as the Electronic Frontier Foundation's (EFF) discussion of Internet privacy. There were also contests like "Capture the Flag" (in which teams of hackers competed to play a "digital" version of the game), "Leetest Link" (a computer security version of "The Weakest Link"), "Wardrive" (during which teams drove around Las Vegas trying to pick up as many wireless networks as they could find), and the "2nd Annual DEFCON WiFi Shootout" (during which teams competed to see who could get the greatest possible connection between two 802.11b stations). This article summarizes three of the lectures: "The Insecure Workstation" by Deral Heiland, "Morph" by Kathy Wang, and "Tor" by Roger Dingledine, which the author found to be among the most interesting of those that he attended (as all lectures attended by the author were interesting in one way or another).

The insecure workstation

Many enterprises take varying steps to keep their employees focused and productive. This is not a bad thing, as without productive employees many businesses would fail because no one would accomplish anything. Beyond the everyday proxies that limit what employees can access on the internet, keystroke loggers that record every word you type, and video cameras that record your every move, many enterprises further restrict what functions can be performed on the system level. A common way for System Administrators to lock down employee workstations in an enterprise is to use Windows Group Policy. Various function, such as; Internet Explorer (IE), Windows Explorer, and Command prompt (cmd.exe), can be "removed" from user level access theoretically keeping users from becoming distracted.

Deral Heiland delivered a very informative talk on how to "break out" of the normal restrictions put on workstations through the use of Group Policies in Microsoft Windows. Group Policies are one way that administrators can centralize control of users and groups. Group Policies include options for registry-based policy settings, security settings, software installation, scripts, folder redirection, Remote Installation Services, and Internet Explorer maintenance. As Mr. Heiland pointed out in his lecture, and reiterated here, this is not an attack on Microsoft but merely a demonstration of mistakes that administrators can make in trying to enforce Group Policies.

Why would administrators prevent users from performing certain functions on a system? Some reasons are—

- To prevent users from "messing up" the system
- To prevent users from accessing privileged data or applications
- To stop malicious users

Typically, Group Policies are used to implement restrictions on users. However, there are some common and dangerous misconceptions often held by administrators—

- "I couldn't get around the restrictions, so they must be fine."
- "The end users aren't smart enough to figure a way around the restrictions."
- "If they are able to get around the restrictions, they will not be able to do anything anyway."

Exploiting group policies

Basically, nothing extraordinary is needed to "break out" of the insecure workstation. You need only to use what is available; i.e., IE, notepad, help screens, or command line. These techniques work on Windows 95 through Windows 2003 and most (if not all) of the versions in between. Mr. Heiland provided these basic exploits—

- Help screens
- Loop back 127.0.0.1 to shares
- USB memory drives
- Older versions of applications
- Trigger errors (i.e., Dr. Watson)
- Security alert pop-ups
- Non-associated extensions

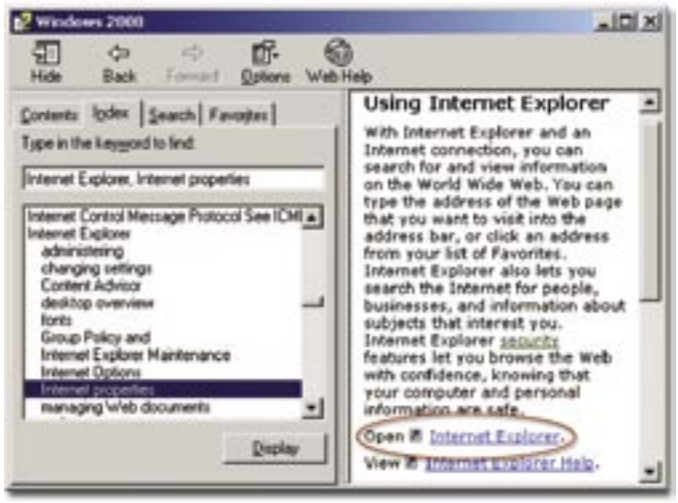

Figure 1: Using Windows Help to access IE

For example, if a workstation is restricted and does not allow access to IE, but the user can access Windows Help, then the user can easily get to IE.

In Figure 1, searching for "Internet Explorer" in Windows Help brings up a link to open IE. Once IE is open, there is a Windows Shell exploit that can be used to open various applications and browse the file structure, including gaining access to a USB drive that could contain other exploit tools, which could come in handy to a penetration tester or even a malicious user. For as in any penetration test (or during a malicious penetration), having the right tools available at the right time can allow a tester (or hacker) further access along the network.

Morph

One of the steps in performing a penetration test (or a malicious penetration) is to perform a footprint analysis of the target system(s). Incorporated in this phase might be a port scan to determine open ports and services on a target host. Probably the most well-known port scanner is a tool called Nmap ("Network Mapper"). Nmap allows a user to remotely map the open ports on a system. Not only does Nmap resolve open ports, but it is also capable of performing Operating System (OS) identification based on the response of the target system. This is where Morph comes in. Morph is a tool that allows a system running one OS [as of this review, it is currently available for Linux and is in development for Open Berkeley Software Distribution (OpenBSD), FreeBSD, and NetBSD] to emulate another OS, in essence fooling any network mappers into associating the wrong OS with the systems they are scanning.

Morph is built on Packet Purgatory, which is described as being—

> *"…a userland library that provides an Application Programming Interface (API) to a "network wedge" that can sit between a system's Internet Protocol (IP) stack and the wire. Outbound packets may be modified after a kernel is done with them but before they have been sent out on a Network Interface Card (NIC), and inbound packets may be modified after they have been received by the NIC but before the firewall is aware of them."*

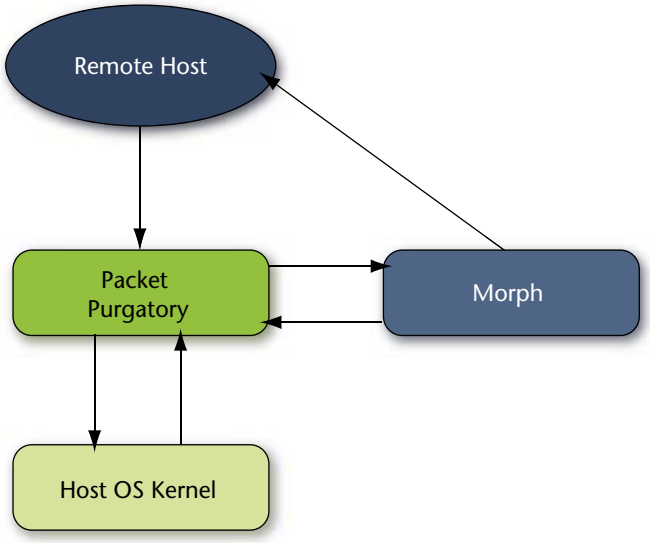

Figure 2: High-level Morph architecture

Morph is able to defeat OS scanners such as QueSO, Nmap, and Xprobe/Xprobe2. Currently, it is under development to defeat OS scanners p0f and RING/Snacktime. These latter scanners use passive fingerprinting to determine the target OS—basically relying on correlating the timing of responses and their respective OSs. Morph relies on a state table to maintain session sequence number offset information. It can then modify the packets to emulate a different OS than that which is actually running.

Table 1 shows how Morph handles different packets sent to it by the originator. Whether the port is open or closed will determine how Morph will respond.

This tool is very useful. Being able to emulate a different operating system than what is actually being used allows an entity to effectively defend itself by going on the offensive and not allowing "hackers" the opportunity to focus their attacks—all the while remaining private.

The Onion Routing (TOR)—Facilitating anonymous communications for Transmission Control Protocol (TCP)

pri•va•cy *n.*: freedom from unauthorized intrusion : state of being let alone and able to keep certain esp. personal matters to oneself

Privacy is a major concern for many Internet users. From checking E-mail, to surfing the Web, to online banking transactions, users may be concerned about who is watching their transactions. Encryption may secure data in transit, but it does not hide the path the data takes. Being able to perform these day-to-day operations with both security and privacy—knowing that no one can monitor their communications and keeping their information private and secure—would be ideal for users that wish to perform transactions in secret. This type of communication could especially come in handy for out armed forces in order to further secure communications—which may be the reason the Navy is sponsoring this type of

		Morph Handling Status		
	Nmap Packet Types	**Inbound**	**State Table**	**Outbound**
Open Port	TCP Sequence Test	Pass packet to OS	Add SYN connection	Send response packet to reflect emulated OS
	SYN with Options	Pass packet to OS	Add SYN connection	Send response packet to reflect emulated OS
	NULL with Options	Respond on behalf of emulated OS	Don't care	Don't care
	SYN-FIN-URG-PSH with Options	If OS accepts it, pass to OS. Otherwise, respond on behalf of emulated OS	Add connection	If applicable, send response to reflect emulated OS
	ACK with Options	If connection exists, pass packet to Os. Otherwise, respond on behalf of emulated OS	If part of existing connection, add ACK connection	Send response packet to reflect emulated OS if part of existing connection
Closed Port	SYN with Options	Respond on behalf of emulated OS	Don't care	Don't care
	ACK with Options	Respond on behalf of emulated OS	Don't care	Don't care
	PSH-FIN-URG with Options	Respond on behalf of emulated OS	Don't care	Don't care
	UDP Packet	Respond on behalf of emulated OS	Don't care	Don't care

Table 1: Morph response to Nmap 3.50

application. TOR, or The Onion Routing , is an attempt to solve this problem by bringing onion routing to the masses through open-source distribution.

Onion Routing is a distributed overlay network designed to render TCP-based applications such as Web browsing, secure shell, and Instant Messaging anonymous. It is a flexible communications infrastructure that is resistant to both eavesdropping and traffic analysis. Onion Routing accomplishes this goal by separating identification from routing. Connections are always anonymous, although communication need not be. Communication may be made anonymous by removing identifying information from the data stream. Onion Routing can be used by a variety of unmodified Internet applications by means of proxies (a non-invasive procedure) or by modifying the network protocol stack on a machine to be connected to the network (a moderate or highly-invasive procedure).

TOR was presented at DEFCON by one of its authors, Roger Dingledine, and was described as "the next generation of Onion Routing." TOR is a project that focuses on the anonymity of the connection (pipe), not on what type of communication is conducted (i.e., Hypertext Transfer Protocol [HTTP], Secure Shell [SSH], or File Transfer Protocol [FTP]), thereby combining the previously used approaches of a mix and a proxy.

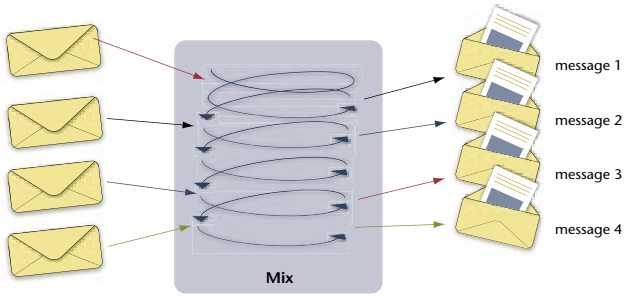

Figure 3: Conceptual diagram of a "Mix"

Basically, the mix receives communications, encrypts them, and forwards them to the next destination in random order, thereby making each network conversation difficult to track.

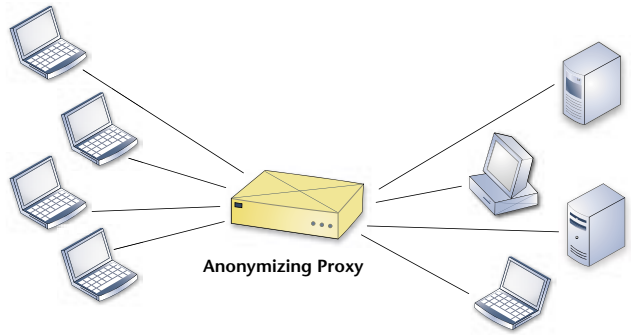

Figure 4: Conceptual diagram of a Proxy

A proxy works by taking a request from an originating system, making the request itself, and then forwarding the results to the originator. The requested addressee only sees the proxy as making the request, thereby hiding the originator's address.

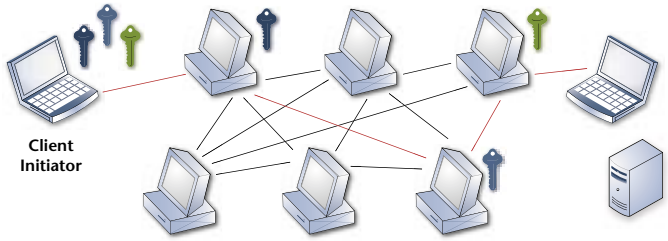

Figure 5: Conceptual diagram of a TOR

TOR creates a "virtual circuit" through a series of Onion Routers, by negotiating keys for encryption along the way. As the TCP traffic is streamed through the circuit, a layer is unwrapped by a symmetric key at each node, which, in turn, reveals the next node.

Closing and future directions

DEFCON is billed as the largest gathering of the hacker underground and appears to increase in size every year, with both technologies and "social" lectures offering interesting information. The Electronic Frontier Foundation presented a two-hour panel discussion on various privacy and First Amendment cases involving the Internet, cases that affect all Americans. Many lecturers provide their Web-site addresses so that attendees can further investigate and even try out what has been presented.

Admission to DEFCON is $80.00 per person and offers an excellent opportunity to learn of new technologies and to meet fellow "hackers." Personnel in Information Technology are urged to attend—DEFCON is probably the best IA education that $80.00 can buy. ■

About the Author

Louis Lerman

Louis Lerman works in the field of IA performing penetration tests and risk assessments for clients such as Fortune 100 firms and various U.S. Government agencies. He holds a B.A. in American Studies (1997) and a B.S.B.A. in Computer & Information Sciences (2001), both from the University of Florida. He has attended DEFCON twice, in 2002 and in 2004, and plans on attending future installments of it.

Evidence-based Health Care and Information Assurance (IA)

by Dr. Laurence Loeb

Information assurance (IA) has historically concerned itself with the collection, transmission, and reception of valid information. In general, the standards used in establishing informational validity are dependent on the specific kind information that is being considered. Health care information has historically been in a grey area as to validation, due to the multiplicity of potential information sources and concerns over patient confidentiality (recently codified by the Health Insurance Portability and Accountability Act (HIPPA) regulations). This has proven problematical to those trying to make informed decisions regarding health care policies and outcomes for a particular patient population, like those found in the military. While military health care has exemption from privacy concerns when compared to the civilian population, relevant medical information may still repose in many places within the health care system. The IA standard of validity when applied to medical information requires new thinking about how this kind of information is made available to those that need it.

Evidence-based medicine

One possible solution is the paradigm of evidence-based medicine that has arisen in the last decade. One working definition of evidence-based medicine was developed by the Cochrane Collaboration—so named for the Oxford, England-based Cochrane Centre. This collaboration included representatives from nine countries, and was aiming to facilitate randomized trails in several areas of health care.

Chalmers describes evidence-based medicine thus—

"Evidence-based medicine is the conscientious, explicit, and judicious use of current best evidence in making decisions about care of individual patients. The practice of evidence-based medicine means integrating individual clinical expertise with the best available external clinical evidence from systematic research. …Increased expertise is reflected in many ways, but especially in more effective and efficient diagnosis and in more thoughtful identification and compassionate use of individual patients' predicaments, rights, and preferences in making clinical decisions about their care. By best available external clinical evidence we mean clinically relevant research, often from the basic sciences of medicine, but especially patient centered clinical research into the accuracy and precision of diagnostic tests (including the clinical exam), the power of prognostic markers, and the efficacy and safety of therapeutic, rehabilitative, and preventive regimens. External clinical evidence both invalidates previously accepted diagnostic treatments and replaces them with new ones that are more powerful, accurate, efficacious and safe."

Assuring informational flow

Thus, evidenced-based medicine demands a continuous flow of current and valid external information that is relevant to the individual's medical situation. Automated information retrieval systems that data-mine medical information sources would seem to be a necessity in attempting this kind of medicine for no other reason that it is impossible for any single practitioner to read all of the currently available medical informational sources. It is a challenge to the IA professional involved in the evidence-based effort to ensure the timely and error-free delivery of the informational content so necessary for the process to succeed.

The Agency for Healthcare Research and Quality has several resources available that can help the IA professional to deliver relevant medical information. Their Web page at http://www.ahrq.gov/clinic/epcix.htm lists various evidence-based reports that would be useful to the clinical practitioner. By using informational resources like this one, IA validity concerns about the content supplied to medical professionals can be directly addressed.

Also, the Department of Health and Human Services (HHS) has recognized the informational needs of health care, and committed itself in July of 2004 to a 10-year plan that builds a national electronic health information infrastructure in the United States. It has as its centerpiece the interoperable Electronic Health Record (EHR), which would allow clinicians access to needed information wherever treatment is rendered.

As an example of a Federal health information technology program that has provided useful information, the Veterans Administration (VA) provides to physicians, registered nurses, dentists, optometrists, podiatrists, nurse anesthetists, physician assistants, and other staff an EHR system known as VistA. The VA's work on the evolution of this EHR and diagnostic imaging is leading the field. The VA first demonstrated the effectiveness of bar coding for improving patient safety in hospital drug administration. This contributed to the FDA's development of regulation requiring bar codes on drug products.

Another example of a federal program that brings needed critical information to health practice is Department of Defense's (DoD's) Pharmacy Data Transaction Service (PDTS), which is linked to the DoD's EHR system. This utilizes a centralized data repository that records information about prescriptions filled for DoD beneficiaries through Military Treatment Facilities (MTFs), the civilian pharmacy network, and the TRICARE Mail Order Pharmacy program. PDTS enhances patient safety and quality of medical care by reducing the likelihood of adverse drug-to-drug interactions, duplicate drugs prescribed to treat the same condition, and the same drug obtained from multiple sources. This system has detected more than 117,000 potential Level 1 drug interactions over the last three years.

A dental example

These concepts have been applied to dentistry, due to the relative lack of complexity that dental diagnosis exhibits compared to general medicine. It is also true that dental services have a demonstrable decades-long history of reporting dental services to payment agencies through procedure codes that are less complex than those of general medicine. The American Public Health Association expressed its views on this in their position statement 9706. In the statement, they—

"support federal agencies such as the Health Resources and Services Administration, National Institute for Dental Research, Agency for Health Care Policy and Research, the Centers for Disease Control and Prevention, the Health Care Financing Administration, the Veterans Administration, as well as state health agencies and the health insurance industry in adequately funding systematic reviews and research projects which provide further evidence of efficiency and cost effectiveness of oral health care."

One company has developed a dental evidence-based system. The attributes of the system were developed by Oral Health International, a dental research and analysis firm, in conjunction with the Center for Health Systems Research and Analysis at the University of Wisconsin, Madison. The Dental Wellness Index (DWI) is a statistically modeled system intended to assess the quality of dental care rendered to recipients in terms of clinical quality, clinical effectiveness, efficient and treatment outcome measurement scores. The DWI consists of a data collection form, the attributes of the DWI (issues, components, weights, measures, and utilities), a series of algorithms that interpret the data collected and a custom software program used for interpretation.

The DWI assesses the appropriateness of treatments, prioritizes treatment in order of merit, assesses treatment outcomes, and monitors cost effectiveness and cost efficiency. It takes a specific patient's information and compares it to relevant historical data that has been collected, providing a comparison that can be measured. Changes in the index (positive or negative) can also provide an assessment of the treatment rendered. In this way, the DWI provides a validation of treatment that is expressed in procedure codes. The information that is provided by the procedure codes themselves when presented in a linear list (such as can be found in a patient's treatment record or in an insurance submittal form) is less than when those same codes are analyzed into the DWI. The DWI is a form of metadata on the health care procedures, providing a handle onto the entire course of treatment rather than just the snapshot of what was done on one particular day.

Summary

The problem of assuring a valid informational flow in medicine and dentistry is non-trivial. It requires both excellent sources of information, peer review of those sources, and the ability to apply that information to the problem at hand. The DoD, as well as the VA, have made significant attempts to address this problem, with some operational successes.

For example, DoD has transferred clinical information on over 2.27 million prior service members to the VA by use of the Federal Health Information Exchange (FHIE). The data comes from the DoD Composite Healthcare System and is stored in the FHIE Data Repository. VA clinicians can view it there.

With the current emphasis (and funding) on health informational technology coming from the HHS, the DoD is tasked to provide its expertise to other branches of the government. HHS, DoD and the VA have endorsed 20 sets of standards to make it easier for information to be shared across agencies and to serve as a model for the private sector. This kind of effort must be encouraged across the board to allow for informational sharing among all the stakeholders in the health care field. Only when all of the parties can understand each other can the promise of deriving relevant information for healthcare be realized. ■

About the Author

Dr. Laurence Loeb

Dr. Laurence Loeb is a Principal at pbc enterprises. Dr. Loeb has authored "SET: Introduction and Technical Reference" (Artech House, 1998) as well as "Hackproofing XML" (Syngress, 2003). He may be reached by E-mail at larryloeb@prodigy.net or via phone/fax at 203/281-6674.

continued from page 9...

"Preventing Widespread Malicious Code Infections"

users receive copies of the same worm via E-mail. No matter how obvious the malicious nature of the worm may be, some users will accidentally run it, no matter how much education and training they have received. People make mistakes! Of course, a worm that is not so obvious is likely to be run by more users, and a worm that appears to be a legitimate message from an authority could fool many users.

Future trends in malicious code

Based on the increasing threats posed by widespread malicious code, it seems almost certain that the worst is yet to come. Several years ago, it took months on average for malicious code to be released after a new vulnerability was announced. This gave organizations the opportunity to eliminate vulnerabilities before they could be exploited. Today, worms are often released within weeks—and in the case of worms such as Witty, within a few days. Some attacks use so-called zero day exploits, meaning that the malicious code is released before the exploited vulnerability has even been publicly announced. In such cases, organizations will need to rely on security controls that prevent the vulnerabilities from being exploited, rather than focusing on vulnerability elimination. Other likely trends involving widespread malicious code infections are as follows—

- **Increased host and domain targeting**—Several recent worms have been configured to launch a denial of service attack from infected machines against a particular host or domain on a particular date. A few recent worms have contained domain-specific programming, such as not spreading to certain domains or harvesting E-mail addresses for only a certain domain.
- **Widespread infections through Web sites**—In 2001, the *Nimda* worm infected Web pages so that many hosts that visited those Web pages were also infected. In June 2004, attackers compromised Web servers belonging to several organizations and placed malicious code on Web pages that then attempted to place Trojan horses on Web browsing hosts. Although there have been few widespread incidents involving infected Web sites, it certainly seems likely that at some point, malicious code will be released that infects many Web pages in a short time, leading to the rapid infection of many Web browsing systems.

Conclusion

Widespread malicious code infections pose an increasingly severe threat to networks and systems. Organizations need to use several types of controls to prevent infections, because no single control can prevent all infections. This article has presented some of the most effective and commonly used controls, but there are many other controls that may also be valuable in mitigating threats. Although mass mailing and network service worms are currently the most prevalent threats, organizations should also carefully consider other types of malicious code and other transmission mechanisms when establishing preventative measures. ■

About the Author

Karen Kent

Karen Kent holds a B.S. in Computer Science from the University of Wisconsin-Parkside and an M.S. in Computer Science from the University of Idaho. Ms. Kent co-authored NIST Special Publication (SP) 800–61: Computer Security Incident Handling Guide and SP 800–68: Guidance for Securing Microsoft Windows XP Systems for IT Professionals: A NIST Security Configuration Checklist. She has also contributed to several books on information security and published over a dozen articles on intrusion detection. She may be reached at kent_karen@bah.com.

letters to the director

A reader of *IAnewsletter* Volume 7 Number 2 Fall 2004 sent IATAC a note regarding the content of our publication. The following summarizes the reader's comments and our response:

I'm...wondering if I'm working in the right office. Admittedly I am no technical guru of any kind. I however have been in the Information Assurance business for over 3 years now. I feel that I have enough experience in IA to know relevant from irrelevant and after reading through your latest issue I find that most of the information contained within it is irrelevant to me as an IA professional. I mean, who is this really speaking to..."

The editorial staff of the IAnewsletter spends considerable time deciding what we should publish in our magazine. Since IATAC's primary focus and mission is to, "provide the DoD a central point of access for information on Information Assurance emerging technologies in system vulnerabilities, research and development, models, and analysis to support the development and implementation of effective defense against Information Warfare attacks," it would follow that our premier publication would mirror the mission and maintain a focus that would appeal to that community. IATAC operates under the direction of a Government Steering Committee whose backgrounds are primarily Research and Development and academia, however, we strive to ensure we are relevant to all our readers. We spend considerable time searching for, selecting, and editing topics that would be appropriate for an R&D or academic environment. At the same time I know that we must grow and reach out to other communities that IA impacts. I agree that the past issue was technical in nature but I hope you find previous (and future) editions less technical. We try to offer technical, practical, managerial and useful information for military leaders and operational level users. As a closing comment one of our senior DoD leaders wrote that the issue you comment on was our best edition ever.

If you have comments or suggestion please forward them to us at iatac@dtic.mil as we welcome you thoughts.

what's new

Intrusion Detection Systems Tools Report

The Intrusion Detection Systems (IDS) Tools Report is the latest of IATAC's updated tools reports. This report is separated into two distinct sections—the first being general information on IDS. This section provides a brief explanation of intrusion detection and prevention systems and the necessity of them. In this report, an Intrusion Detection System (IDS) is described as a system which attempts to detect intrusion into a computer or network by way of observation or audit. Intrusion Prevention Systems (IPS) go one step further and not only detect the attacks but attempt to prevent them from succeeding as well. The second part of the report contains an index, with just some of the various IDS/IPS tools available today. The index is divided into three sections—host-based, network-based, and suites of intrusion detection/prevention systems tools.

For instructions on obtaining a copy of the report, either visit us on line at http://iac.dtic.mil/iatac/products.html or refer to the IATAC Product Order form located on the back page of the IAnewsletter.

continued from page 23…

"The Importance of High Quality Information Assurance (IA) Metrics"

work defense is also supportive of strategic goal one—protect information. An activity that doesn't map to one of the five major IA goals may not adequately support the DoD IA program. Similarly, an activity that strongly maps to multiple IA strategic goals is arguably more relevant and valuable. Additionally, mapping activities to the DoD IA strategy is often an aid in selecting appropriate performance objectives.

2. **Quantitative**—IA metrics should yield quantitative rather than qualitative information. This does not suggest that qualitative inputs cannot be considered, but that qualitative inputs such as opinion surveys should be aggregated and presented as a numeric value. An example of this would be percentage change in satisfactory versus unsatisfactory ratings. Numerical metrics make a more compelling argument to support executive-level risk and resource decisions.

3. **Reasonable**—Data should be immediately obtainable or easily collected through a survey or from a data repository. In other words, the value of the data collected should not exceed the cost of collection. IA metric collection is not the ultimate purpose of most programs—they are a means, rather than an end. The cost of measuring should not impinge on mission accomplishment. It should be noted here that externally driven IA metrics often result in additional costs beyond their immediate value to the program being measured. However, the reasonableness criteria for such IA metrics are the responsibility of the requester.

4. **Verifiable**—IA metrics should be evidence-based and that evidence should be objectively verifiable. Third party reviewers or auditors should be able to assess data and concur with a metric's result. This standard excludes the opinion-based questions discussed early. Why require such rigor? Remember, IA metrics results will be used to justify decisions and resource requests. Basing metrics on data rather than opinion reduces the risk that decisions will be based on inadequate or inaccurate data.

5. **Trendable**—Organizations should not create new metrics every quarter or even every year. IA metrics should ask questions that will be of interest for an extended period of time. Such metrics substantiate the results of business or financial decisions—"yes, that investment in cryptography enhanced the security of data"—"no, adding resources to program Y did not increase efficiency." Trendable metrics can still be flexible—the performance target of a metric can adjust from year to year in response to environmental and resource changes. An IA metric is typically "dropped" from consideration after an organization has attained a desired level of performance for an extended period of time.

6. **Useful**—Results should yield information that is important in mission or financial decision-making. Usefulness may not be transitive—an IA metric that is useful to one organization may not be useful to the next. The manager of an IA training program may not be directly interested in metrics about Common Access Card program implementation. Similarly, managers at different levels within an organization have differing requirements for decision support information. The quality metrics of an IA activity being managed at lower levels of the organization would not necessarily answer the need of management at the Departmental level, and visa versa. That is not to say that there may not be common metrics. Some local IA metrics that can be rolled up through the organization and be very meaningful at the highest levels of management. But these are the exception, not the rule.

7. **Indivisible**—Data should be collected at the most discrete, unanalyzed level possible. It may be important to collect the number of accredited systems, but that number in the context of the total number of systems is more useful. Discrete measures and metrics may be reused in unexpected ways by other organizations. Also, there are cases where data of disparate origin can be combined or otherwise correlated to provide a quality IA metric. For example, turnover rates collected in support of human resources programs can be combined with IA metrics to predict potential resource short falls.

8. **Well-Defined**—IA metrics and their attributes must be well defined. Metric creators should document the characteristics of their metrics—What is the frequency of a metric? What formula is used to calculate a metric? What elements or evidence are required to substantiate a metric? What might the metric indicate if the result over time trends up or down?

To ensure that quality IA metrics are collected, DoD and component organizations should promulgate IA metrics standards and manuals. Documents that institutionalize the processes for selecting, analyzing, and reporting on IA metrics would also add consistency to DoD's compliance reporting and could advance the goals of the Net-Centric Data Strategy. These standards and manuals should also address the optimal design of data collection surveys. Common tools can also support quality metric identification, collection, analysis, and reporting. IA metrics should be defined using a highly descriptive common template so that the users and generators of data specifically understand the proposed interpretation of data. IA data collection and analysis can often be automated. Automating data collection is usually less intrusive and generally more accurate than manual methods of collection.

DoD' senior IA leadership is initiating an enterprise-wide metrics program to improve insight into the IA posture of DoD. As we increase our reliance on IA metrics to support risk and budget decisions across the Department, we must continue to improve IA metrics quality. The eight criteria for quality metrics will help DoD ensure quality, and better assure DoD information and information technology resources. ■

About the Authors

Vivian Cocca, CISSP, OSD/NII IA

Ms. Cocca is assigned to the Office of the Assistant Secretary of Defense (Networks and Information Integration). She currently serves as the Senior IA Analyst for Strategy, Plans, and Outreach for the Director, Information Assurance (IA), office of the Deputy Chief Information Officer (DCIO). Ms. Cocca has 20 years experience with the Department of Defense and for the last seven, served on the staff for the Secretary of Defense. She has extensive experience in areas of strategic analysis, intelligence methods, and management and program oversight. She earned a Bachelor of Arts Degree in Geography from the University of California, Santa Barbara and a Masters Degree in Strategic Intelligence from the Defense Intelligence College. She also earned her CIO certification from the Information Resources Management College, and is a Certified Information Systems Security Professional (CISSP). Past assignments include the Pentagon, U.S. Army Intelligence, Special Operations, Fort Bliss, and Fort Huachuca.

Steven Skolochenko, CISSP, CISA, CISM

Mr. Skolochenko is an Associate with Booz Allen Hamilton who has extensive federal experience in information security, including extended efforts at Treasury, Justice, the U.S. Postal Service, and the U.S. Army. He holds masters degrees from both the George Washington University and the American University in Management of Science and Technology and Management Information Systems respectively. Mr. Skolochenko is a frequent speaker at information systems security conferences and served as a member of the steering committee of the Federal Computer Security Program Managers Forum.

Jonathan Smith, CISSP

Mr. Smith is an Associate at Booz Allen Hamilton who specializes in policy, program development, strategic planning, and implementation of information assurance and critical infrastructure protection for U.S. government agencies. He has seven years of IT and IA experience and has supported a wide range of federal agencies. Mr. Smith is also co-author of NIST Special Publication 800-47, "Security Guide for Interconnecting Information Technology Systems." Mr. Smith's background also includes technology training, federal government service, network engineering, and a degree in Philosophy.

References:

Security Metrics Guide for Information Technology Systems. National Institute of Standards and Technology—Special Publication 800-55.

International Standard ISO/IEC 21827: 2002–10–10, Information Technology—Systems Security Engineering - Capability Maturity Model (SSE–CMM®).

OMB Circular A–11, Preparation, Submission and Execution of the Budget (Revised 07/16/2004).

DoD Net-Centric Data Strategy. DoD CIO. May 9, 2003.

DoD Information Assurance Strategic Plan.

IATAC Conference, Meeting, & Event Planning

Providing technical and administrative support for scientific, technical, and DoD-related information assurance management conferences, symposia, workshops, and other meetings. We will coordinate all resources to ensure your event is a success!

Services

Pre-Event Support

- ❏ Site selection
- ❏ Catering arrangements
- ❏ Contract negotiation
- ❏ Promotion and marketing
- ❏ Event support materials
  - – Agenda
  - – Notebooks and folders
  - – Presentation materials
- ❏ Security and registration

On-Site Support

- ❏ Coordination with caterers
- ❏ Check-in of registrants
- ❏ Document control
- ❏ Security problem resolution (if required)

Post-Event Support

- ❏ Create and assemble event proceedings
  - – CD-ROMs
  - – Hard copies
- ❏ Distribute event proceedings
- ❏ Generate final report

Benefits

A Proven Approach

- ❏ Detailed pre-planning expertise
- ❏ History of numerous successfully planned and executed events
- ❏ Expertise in policy adherence for conducting classified conferences
- ❏ Commitment to sponsor needs

Hotel/Sales/Catering

We work closely with hotels to block rooms and negotiate a predetermined conference room rate, coordinate food and beverages for breaks, lunches, and receptions.

Attendees

We work closely with each attendee to ensure we have all the appropriate registration information, security forms, and fees.

Event Marketing

We identify and take advantage of all appropriate promotional and marketing opportunities with professional associations, newsletters, other periodicals, and Web sites.

IATAC possesses world-class telecommunication, graphics, printing, and reproduction capabilities, providing service and support to guarantee the highest quality conference preparation materials, brochures, posters, presentations, proceedings, and product displays, both electronic and hard copy. IATAC also possesses outstanding multimedia presentation capabilities, which includes Web page development and on-line registration.

Classified Session Facilities

We coordinate with the appropriate personnel, ensuring compliance to classified event procedures. We work closely with security personnel and to develop appropriate mailing and storage instructions for classified presentations.

Please contact us at the information below.

3190 Fairview Park Drive, Falls Church, VA 22042

Commercial: 703/289-5454
Fax: 703/289-5467
E-mail: iatac@dtic.mil
URL: http://iac.dtic.mil/iatac

Conference/Event Planner
April Perera
703/289-5699
iatac@dtic.mil

Promotional Director
Christina P. McNemar
703/289-5464
iatac@dtic.mil

product order form

Instructions: All IATAC **LIMITED DISTRIBUTION** reports are distributed through DTIC. If you are not a registered DTIC user, you must do so **prior** to ordering any IATAC products (unless you are DoD or Government personnel). **To register On-line:** http://www.dtic.mil/dtic/registration.
The *IAnewsletter* is **UNLIMITED DISTRIBUTION** and may be requested directly from IATAC.

Name _____ DTIC User Code_____

Organization _____ Ofc. Symbol _____

Address_____ Phone _____

_____ E-mail _____

_____ Fax _____

Please check one: ❑ USA ❑ USMC ❑ USN ❑ USAF ❑ DoD
❑ Industry ❑ Academia ❑ Gov't ❑ Other

Please list the Government Program(s)/Project(s) that the product(s) will be used to support: _____

_____

LIMITED DISTRIBUTION

IA Tools Reports (softcopy only)

❑ Firewalls ❑ Intrusion Detection ❑ Vulnerability Analysis

Critical Review and Technology Assessment (CR/TA) Reports

❑ Biometrics (soft copy only) ❑ Computer Forensics* (soft copy only) ❑ Configuration Management
❑ Defense in Depth (soft copy only) ❑ Data Mining ❑ Exploring Biotechnology
❑ IA Metrics (soft copy only) ❑ Network Centric Warfare
❑ Wireless Wide Area Network (WWAN) Security

State-of-the-Art Reports (SOARs)

❑ Data Embedding for IA (soft copy only) ❑ IO/IA Visualization Technologies
❑ Modeling & Simulation for IA ❑ Malicious Code

* You MUST supply your DTIC user code before these reports will be shipped to you.

UNLIMITED DISTRIBUTION

Hardcopy *IAnewsletters* available

Volumes 4		❑ No. 2	❑ No. 3	❑ No. 4
Volumes 5	❑ No. 1	❑ No. 2	❑ No. 3	❑ No. 4
Volumes 6	❑ No. 1	❑ No. 2	❑ No. 3	❑ No. 4
Volumes 7	❑ No. 1	❑ No. 2	❑ No. 3	

Softcopy *IAnewsletters* back issues are available for download at http://iac.dtic.mil/iatac/IA_newsletter.html

Fax completed form to IATAC at 703/289-5467

January

GOV–CON05
January 10–12, 2005
Anaheim, CA
http://www.ncsi.com/govcon05/overview.shtml

2005 DoD Cyber Crime Conference
January 10–14, 2005
Westin Innisbrook Golf Resort, Clearwater, FL
www.DoDCyberCrime.com

National Reconnaissance Office
January 25–26, 2005
NRO Complex, Chantilly, VA
http://www.fbcinc.com/event.asp?eventid=Q6UJ9A007ZWS

Network Centric Warfare 2005
January 25–27, 2005
Ronald Reagan Building and International Trade Center, Washington, DC
http://www.idga.org/cgi-bin/templates/singlecell.html?topic=222&event=5517

TechNet Orlando 2005
January 25–27, 2005
Radisson University Hotel, Orlando, FL
http://www.afcea-orlando.org/documents/TechNet%20Orlando%202005%20Announcement%20(1).doc

February

West 2005 "Beyond Iraq: How Do We Get Transformation Right?"
February 1–3, 2005
San Diego Convention Center, San Diego, CA
http://www.west2005.org/

9th Annual IA Workshop 2005 GIG IA: Enabling the Global Battlefield
February 7-10 2004
Philadelphia, PA
http://iase.disa.mil/iaws9.html

National Threat Symposium
February 8–9, 2005
Las Vegas, NV
http://www.ioss.gov/conf/nts_west.html

Defending America—SPACECOM 2005
February 8–10, 2005
Broadmoor Hotel International Center, Colorado Springs, CO
http://www.rockymtn-afcea.org/2005/

INTELCON 2005
February 9–10, 2005
Hyatt Regency Crystal City, Arlington, VA
http://www.fbcinc.com/intelcon/

RSA® Conference 2005
February 14–18, 2005
Moscone Center, San Francisco, CA
http://www.rsasecurity.com/conference/conf_portal.html

April

21st National Space Symposium
April 4–7, 2005
The Broadmoor, Colorado Springs, CO
http://www.spacesymposium.org/national05/information/index.cfm

REDTEAM2005 Conference
April 6–28, 2005
Sandia National Laboratories, Kirtland AFB, Albuquerque, NM
REDTEAM2005@sandia.gov

4th Annual PKI R&D Workshop "Multiple Paths to Trust"
April 19–21, 2005
National Institute of Standards, Gaithersburg, MD
http://middleware.internet2.edu/pki05/

DoDIIS Worldwide Conference
April 25–28, 2005
Philadelphia Marriott, Philadelphia, PA
http://www.ncsi.com/b2b/DoDIIS05early.asp

IATAC
Information Assurance Technology Analysis Center
3190 Fairview Park Drive
Falls Church, VA 22042

To change, add, or delete your mailing or E-mail address (soft-copy receipt), please contact us at the address above or call us at: 703/289-5454, fax us at: 703/289-5467, or send us a message at: iatac@dtic.mil