IAnewsletter

The Newsletter for Information Assurance Technology Professionals

Volume 5 Number 1 • Spring 2002

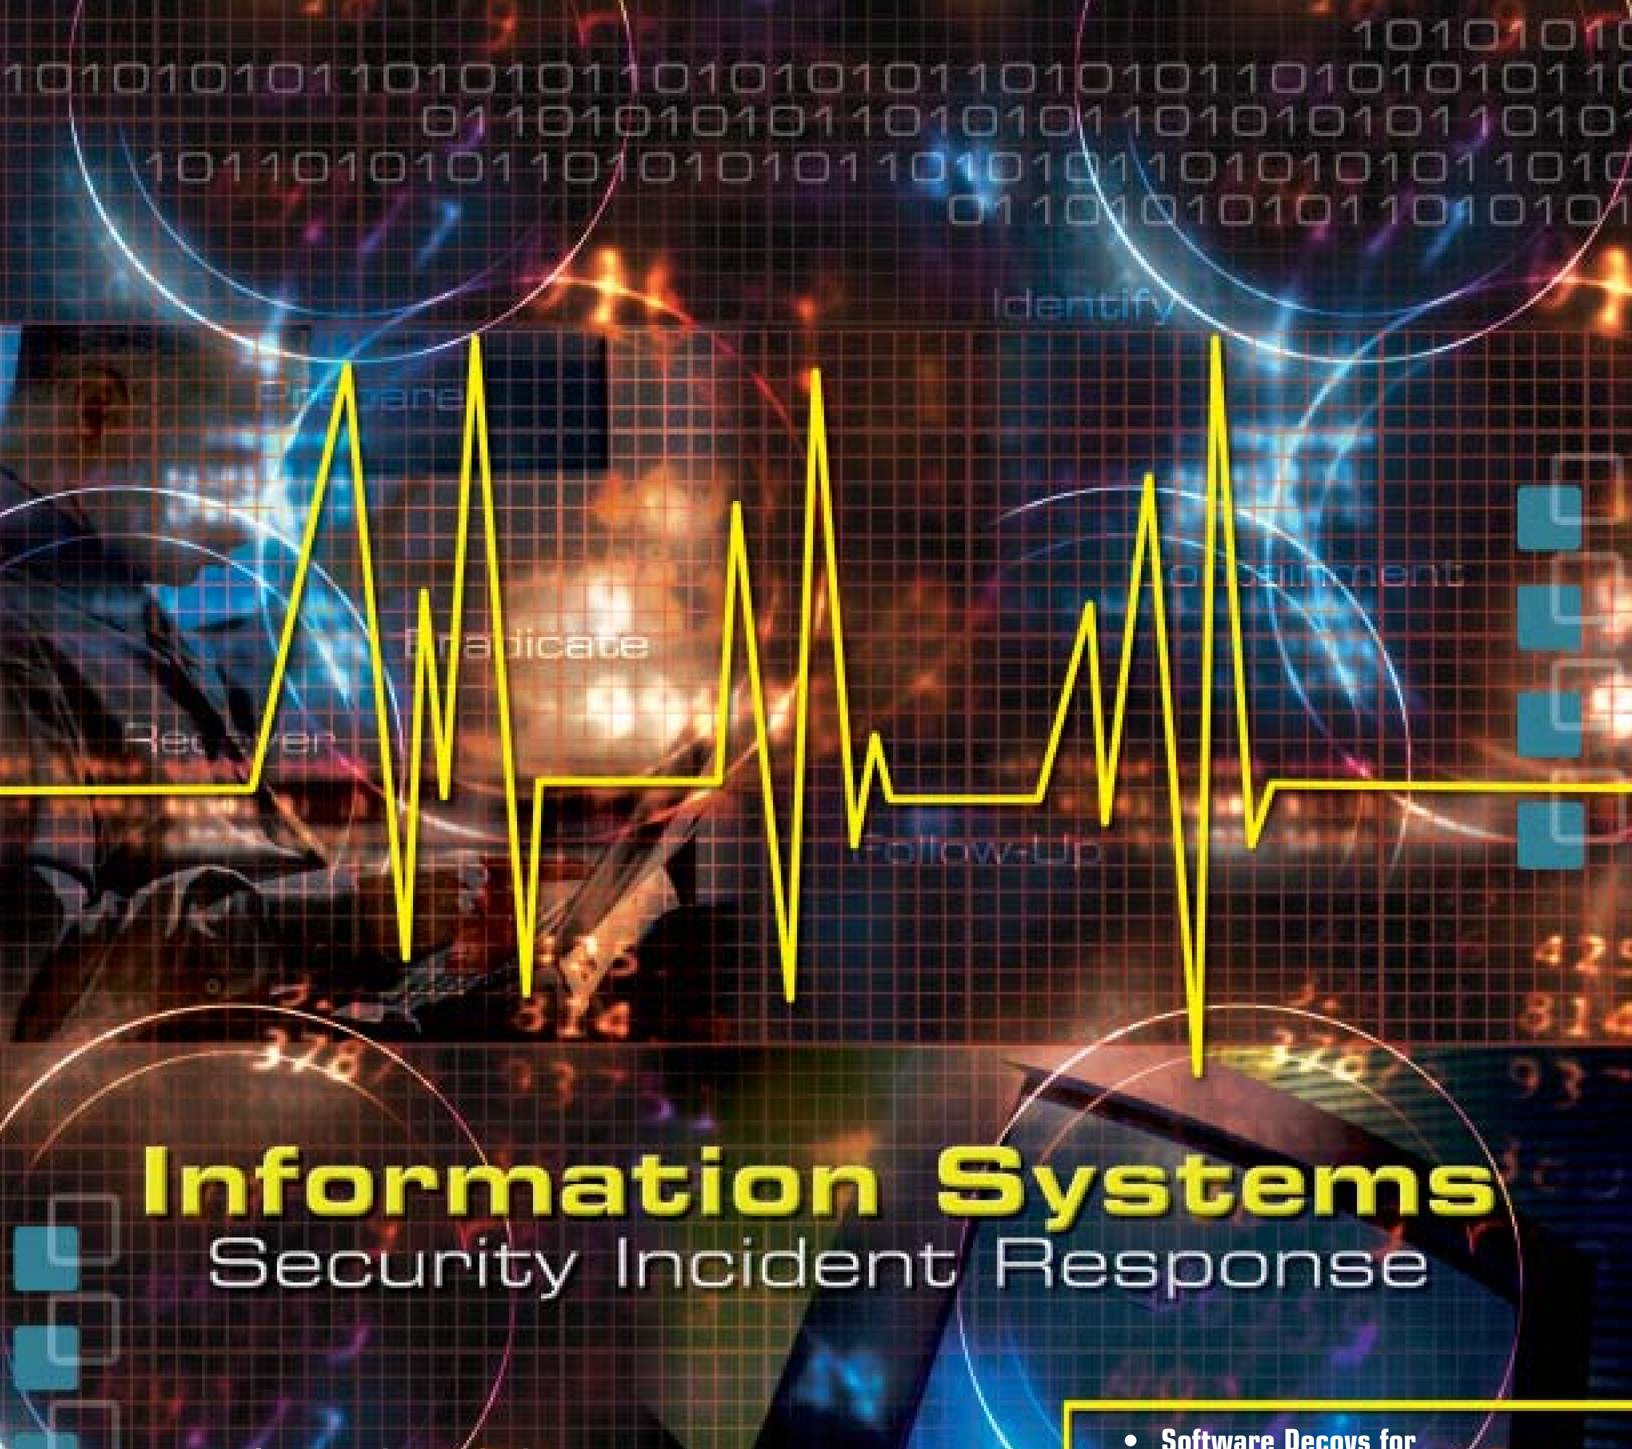

Information Systems Security Incident Response

also inside—

- Software Decoys for Software Counterintelligence
- Phoenix Challenge: Information Operations Concepts and Solutions Exploration in the 21st Century

contents

on the cover

ia initiatives

in each issue

IAnewsletter

Editor
Robert J. Lamb
Creative Director
Christina P. McNemar
Art Director
Ahnie Senft
Illustrations
Maria Candelaria
Information Processing
Abraham T. Usher
Inquiry Services
Peggy O'Connor

IAnewsletter is published quarterly by the Information Assurance Technology Analysis Center (IATAC). IATAC is a DoD sponsored Information Analysis Center, administratively managed by the Defense Technical Information Center (DTIC), Defense Information Systems Agency (DISA).

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

Robert J. Lamb
Director, IATAC
703.289.5454

Submitting Articles

To submit your related articles, photos, notices, feature programs, or ideas for future issues, please request an author's packet from—

IATAC
Christina P. McNemar
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/news_events/author_submission.htm

Article Deadlines

Summer 2002	15 Jun
Fall 2002	16 Aug
Winter 2002/2003	16 Nov

Cover design: Maria Candelaria
Original newsletter design: Christina P. McNemar

Distribution Statement A:
Approved for public release;
distribution is unlimited.

IATAC chat

by Mr. Robert J. Lamb, IATAC Director

I'd like to use this column to introduce some new IATAC members who have joined us over the past couple of months. They collectively and individually bring a wealth of knowledge and experience to IATAC.

Gordon Steele has recently joined IATAC as the Deputy Director. Gordon served in the Marine Corps for 10 years in a variety of assignments including as a Signals Intelligence/Electronic Warfare Officer (SIGINT/EW) and Contracting Officer Technical Representative developing ground-based COMINT/Radio Direction Finding Systems. Gordon's operational experience included service as a SIGINT detachment Operations Officer with the Multinational Force, Beirut. He subsequently entered the commercial sector and has worked 10 additional years in a number of capacities supporting a broad range of the information assurance and information security domain. He has served as a Computer Security Incident Response Manager, prototype development/field support operations manager for a customized data communications surveillance system and an automated information technology planning tool, among many others. Gordon holds an M.S. in electrical engineering.

Ron Ritchey has 16 years of professional experience providing computer and telecommunication services and is an authority in the areas of secure network design and network intrusion. He regularly leads penetration testing efforts where he has had the opportunity to learn first-hand the real-world impact of network vulnerabilities. He is also an active researcher in the field with peer-reviewed publications in the area of automated network security analysis. In addition to his research papers, Mr. Ritchey regularly contributes to security publications including a new book titled *Inside Perimeter Security* which will be published by New Rider in Summer 2002. Mr. Ritchey has authored courses on computer security that have been taught across the country and periodically teaches masters level courses on computer security for George Mason University. Ron holds an M.S. in computer science from George Mason University and is currently pursuing his Ph.D. in Information Technology at their School of Information Technology and Engineering. His doctoral research is attempting to automate network security analysis. Ron is currently working on an IATAC State-of-the-Art Report on Malicious Code.

Rick Aldrich joined IATAC from the Air Force having served over 20 years, the last 15 of those as a Judge Advocate. With a B.S. in Computer Science from the Air Force Academy and a Juris Doctor from UCLA, Rick was uniquely well-versed in the complexities and challenges confronting the Department of Defense within the computer network domain. Adding an LL.M. concentrating in computer law from the University of Houston, Rick went to the Air Force Office of Special Investigations, where he specialized in advising on cybercrime and information warfare-related issues. Since joining IATAC, Rick has addressed a senior-level Information Operations course at the National Defense University on "International Law Issues Related To Crime and Arms Control in Cyberspace." In April he presented at the Army Judge Advocate General's School on the topic, "Domestic Law and Policy in Computer Network Operations." Rick will also be speaking at SANSFIRE 2002 in Boston, MA on the controversial issue, "Do Borders Matter In Cyberspace? Legal Issues Related to the Investigation and Prosecution of Trans-Border Cyber Crimes." Based on his extensive background and insights in this arena, Rick has developed a one-day training course on the domestic, international, and military laws and policies applicable in cyberspace.

IATAC will be participating in a number of upcoming conferences including PACOM's IA Conference and the IEEE Conference on Security and Privacy in May, and the 3rd Annual IEEE IA Conference in June.

Bob

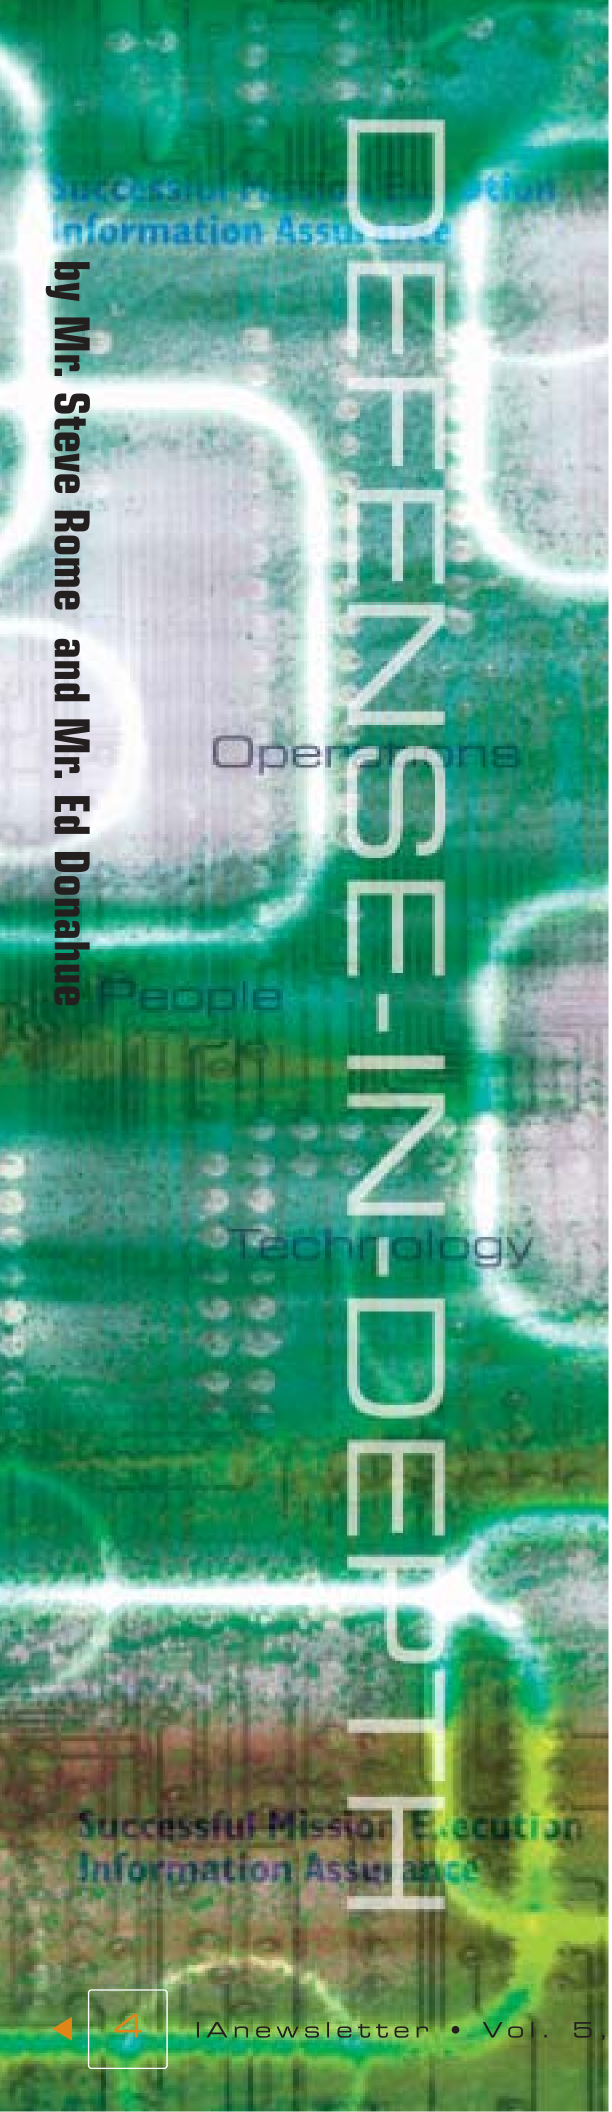

by Mr. Steve Rome and Mr. Ed Donahue

IATF:

At Five Years Old, A Wealth of Knowledge, and Still Growing!

If you have not seen the Information Assurance Technical Framework (IATF) document or attended one of the Information Assurance Technical Framework Forum (IATFF) sessions, you are missing some of the best available opportunities to learn, share, and develop IA technical expertise. The IATF, or Framework, and its accompanying IATFF, or Forum, form a pair. Together, they provide the Department of Defense (DoD) IA community, and now the wider government IA community, and their commercial and academic partners with means of learning from and teaching each other, and of recording the knowledge in a form that practicing IA professionals can use to solve practical problems.

The Formative Years: A Brief History

In 1996, NSA hosted a series of technical sessions relating to network security for several of its customers and partners from across the IA community. These sessions addressed security frameworks ranging from access controls to certificate management. Later that year, the first version of the Network Security Framework (NSF) was published and shared with the community—a few hundred professionals in DoD, DoD vendors, integrators, and Federally Funded Research and Development Center (FFRDC) support personnel. Subsequent versions of the NSF document significantly expanded information on security robustness, security services, security management, and interoperability based on information presented at the Network Security Framework Forum (NSFF). As membership in the forum grew, more information was made available on the NSFF Web site.

By 1998, the almost 2,000 members of the NSFF represented organizations throughout the civil and private sectors, as well as DoD. The information in the NSF was always intended to apply across a wide range of user environments. Because the vast majority of products are employed in a networked environment, in August 1999, NSA changed the names of both the Framework and the supporting Forum to the Information Assurance Technical Framework and the Information Assurance Technical Framework Forum, respectively. This change resulted in version 2.0 of the Framework, which also aligned areas intended to aid users seeking solutions in specific environments (solution frameworks) with the emerging Defense-In-Depth strategy (see Figure 1). By September 2000, the members of the IATFF represented all segments of the IA community. At this point, working with the National Institute of Standards and Technology (NIST), NSA's partner in the U.S. adoption of the International Common Criteria, NSA published version 3.0 of the IATF. This version of the Framework "nationalized" the presentation and content of the document so that it could be adopted by the Federal Government as well as the private sector.

Through Adolescence and Into Maturity

At two inches thick, the IATF document is imposing, but don't be turned off by its size. The IATF may be the most complete compilation of IA security guidance in existence. It offers concise guidance on the full range of information security issues. For beginning information systems security engineers, it is almost a bible. For users faced with securing a system, it addresses many of the security concerns they must address. For some, it is a reference document; for others, it is a text. It is tutorial (vice prescriptive) in nature, in recognition of the fact that many organizations face unique challenges that don't lend themselves to "one size fits all" solutions. The Framework offers insights to improve the community's awareness of trade-offs among available solutions (at a technology, not a product, level) and of the characteristics that are desirable in IA approaches to particular problems. Although the Framework presents a large amount of information, its structure and comprehensive table of contents give readers easy access to topics of interest.

The IATF begins with an explanation of the information infrastructure, its boundaries, the areas of the IA framework, and general classes of threats. It then explains the Defense-in-Depth objectives and elaborates on four Defense-in-Depth technology focus areas—

- Defend the network and infrastructure
- Defend the enclave boundary
- Defend the computing environment
- Supporting infrastructures.

The next chapter discusses the Information Systems Security Engineering (ISSE) and systems engineering processes. An understanding of these processes is helpful in using the IATF. The two processes share common elements: discovering needs, defining system functionality, designing system elements, producing and installing the system, and assessing the effectiveness of the system. Other systems processes—systems acquisition, risk management, certification and accreditation, and life-cycle support processes—are then explained in relation to the ISSE process. These processes provide the basis for the background information, technology assessments, and guidance contained in the remainder of the IATF document.

Chapter 4 of the IATF presents a discussion of the principles for determining appropriate technical security countermeasures. This section includes a detailed description of threats, including attacker motivations, information security services, and appropriate security technologies. Using the methodology described in the ISSE process to assess threats to the information infrastructure allows the identification of vulnerabilities, which is followed by a managed approach to mitigating risks. This section explains how primary security mechanisms, the robustness strategy, interoperability, and Key Management Infrastructure (KMI)/Public Key Infrastructure (PKI) should

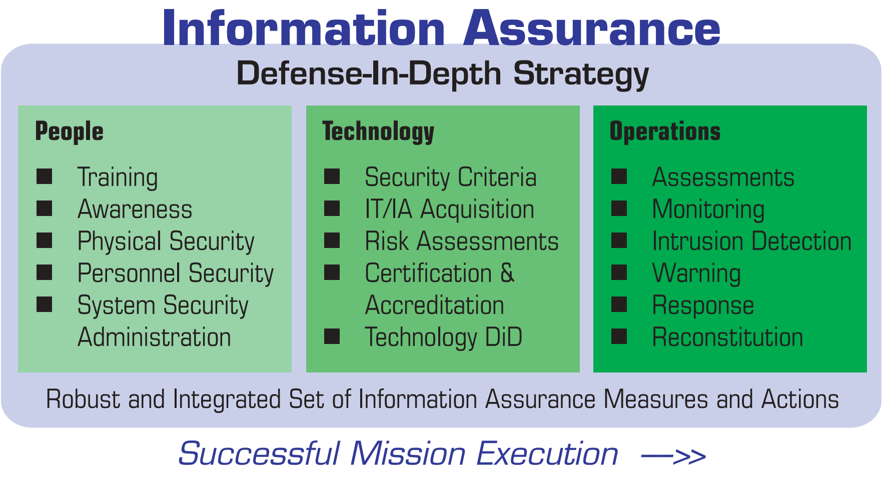

Figure 1. Defense-In-Depth Strategy

be considered in selecting security countermeasures, technology, and mechanisms. These decisions form the basis for the development of appropriate technical countermeasures for the identified threats, based on the value of the information.

The next four chapters deal with the specific technical focus areas of the Defense-in Depth. In each of these four sections, the information follows the same methodology. Essentially, in each of these sections the IATF explains how to identify risks and how to mitigate those risks.

Chapter 5, Defend the Network and Infrastructure, addresses backbone networks and issues with network management. The wireless section of this chapter deals with the special security issues associated with cellular service, pagers, satellite systems, and wireless LANs. The technology assessment section gives guidance on reverse tunneling, virtual private networking, and remote access.

Chapter 6, Defend the Enclave Boundary, deals with control and monitoring of the data flow for external connections to other networks by addressing—

- Firewalls
- Guards
- Remote access
- Virus/malicious code detection
- Intrusion detection
- Multilevel security

This chapter includes 30 pages of firewall guidance, discussing potential attacks and countermeasures for situations in which firewalls might be employed, as well as the role of firewalls in solution sets for defending the enclave boundary.

In Chapter 7, Defend the Computing Environment, the Framework deals with assuring information as it enters, leaves, or resides on clients and servers. In this chapter, one can learn about security-enabled applications, secure operating systems, and host-based monitoring. System administrators also will find a wealth of knowledge to help them better manage their networks. In addition, one section of the chapter examines the technology for secure messaging, secure Web browsing, and file protection. Although the bulk of the Framework addresses technology solutions, some sections of this chapter also deal with the operational aspects of effective network monitoring.

The Supporting Infrastructure chapter (Chapter 8) presents KMI/PKI and detect-and-respond technologies. Here, readers can find more than 100 pages of information on KMI/PKI services and processes and information to help in effectively mitigating the effects of cyber attacks against networks. The detect-and-respond section also includes a discussion of architectural considerations for improving the detect-and-respond posture of an enterprise, evolving paradigms for a detect-and-respond infrastructure, and the technologies available for realizing the processes and functions performed within the secure infrastructure.

Appendix G is intended to share protection profiles; however, today these profiles appear only on the IATF Web site (http://www.iatf.net). In an effort to provide common guidance and a uniform recommendation to its customers for the acquisition of commercial information security products such as firewalls, intrusion detection systems, PKI, and so on, NSA has written many such profiles. These profiles can be used as part of the acquisition specifications. Commercial products can then be tested against these specifications by independent commercial laboratories certified by the National Information Assurance Partnership[1] (NIAP; the U.S. Government's member organization in the International Common Criteria). National Se-

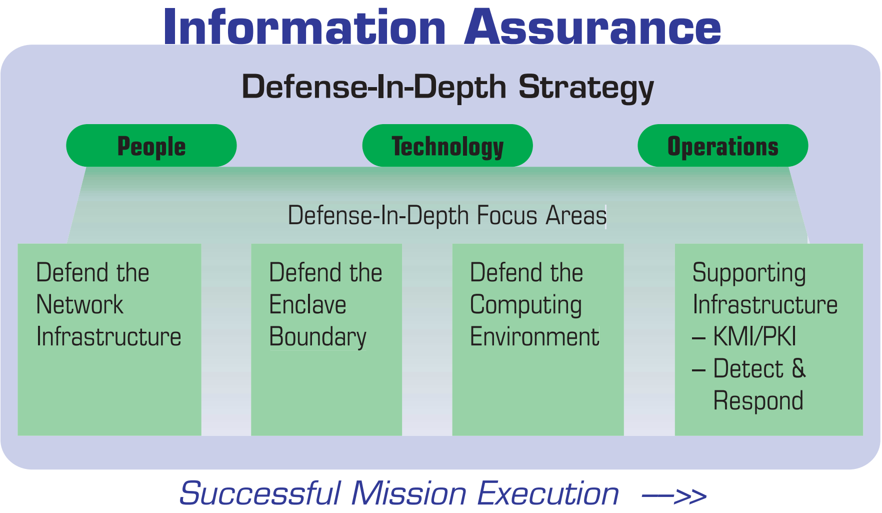
Information Assurance
Defense-In-Depth Strategy
People
Technology
Operations
Defense-In-Depth Focus Areas
Defend the Network Infrastructure
Defend the Enclave Boundary
Defend the Computing Environment
Supporting Infrastructure
– KMI/PKI
– Detect & Respond
Successful Mission Execution —>>

Figure 2

curity Telecommunications and Information Systems Security Policy (NSTISSP) No. 11, issued by the National Security Telecommunications and Information Systems Security Committee (NSTISSC) in January 2000, provided part of the motivation for this program. This policy declares that after July 2002 DoD will purchase only commercial information technology solutions that are NIAP-validated under the internationally recognized Common Criteria (CC) scheme. The Common Criteria, however, does not provide "canned" sets of security requirements; instead, it provides mechanisms for the construction of these sets, from both the user's viewpoint ("protection profiles") and the vendor's viewpoint ("security targets"). The Common Criteria also provides a means of showing that the vendor's product meets the vendor's own security target, and that the vendor's security target—and thus its validated products—satisfies a customer's needs (i.e., conforms to the protection profile the customer wrote or selected as expressing those needs). NSA's protection profile effort is meant to provide a recommended and, it is hoped, uniform set of specifications for these security devices. NSA hopes that this will provide a focus for vendors, and that the vendors will then be motivated to produce products that satisfy customers' requirements as expressed in these protection profiles. Of course, NSA depends on its customers to participate in this process and to ensure that the profiles really express the customers' requirements. Similarly, NSA needs vendor input to ensure that its security requirements are realistic for a commercially marketed product.

A Nurturing Environment is Critical

Today, NSA continues to host IATFF sessions every six weeks. The Forum now has more than 4,700 active members, of which more than 400 regularly attend sessions. These sessions are intended as technical exchanges on various topics related to the IATF. They foster dialogue between developers, users, and IA specialists. This dialogue, in turn, leads to updates of the IATF, as well as providing one of the best-value educational opportunities available in the IA field today. Sessions are free; although, at some venues there is a nominal registration fee to cover expenses specific to that site.

The Forum is supported by the http://www.iatf.net Web site. Here, one can view past topics and presentations as well as summaries of discussions, agendas for future sessions, the IATF and protection profiles written in Common Criteria language, membership registration information, and maps and directions to Forum sessions. If you want to learn more about either the IATF or the Forum, visit the Web site and join today.

Endnote

1. The products can also be tested by CC labs in partner countries and their validation recognized by NIAP under the international recognition program.

Ed Donahue is the Cryptographer for Invicta Networks and also consults for Information Security Systems Inc. in cryptography and information security. He has an M.S. and Ph.D. in Mathematics from Rensselaer. He spent over twenty years at NSA in cryptography, cryptanalysis, and allied fields, including serving as NSA's Chief of Cryptographic Design. He may be reached at 410.461.3574 or edonahue@alumni.bowdoin.edu

Steve Rome is a Senior Associate with Booz Allen Hamilton. He received his B.S. degree in Mechanical Engineering from the University of Maryland, an MBA from Central Michigan University in Management and Supervision, and a M.S. degree in National Resource Strategy from the Industrial College of the Armed Forces, National Defense University. Steve has over 27 years of experience at the National Security Agency in Information Systems Security where he was previously chairman of the IATFF and chief of the IA Solutions Development and Deployment's Architectural Engineering Division.

Phoenix Challenge: Information Operations Concepts and Solutions Exploration in the 21st Century

by Mr. Thomas Sweet and Ms. Sandra Vasile

As the United States Air Force prepares its Expeditionary Air Forces to successfully conduct its Information Operations (IO) missions at home and abroad, it has recently embarked on a grass-roots campaign to identify innovative, state-of-the-art IO technologies and capabilities from within the Department of Defense (DoD), Federal government, industry, and academia that could be used in this campaign. In the face of an austere budget environment, and in response to a myriad of IO challenges and problem sets the Air Force knows it will face in the 21st Century, the Air Force Information Warfare Center, or AFIWC, started a program called Phoenix Challenge, which continues to gain support from within the Air Force and in the IO community.

Phoenix Challenge is an AFIWC and DoD initiative to identify state-of-the-art IO technologies, systems, and capabilities of other Service and DoD IO organizations, the Federal government, industry, and academia that could be leveraged and incorporated into current and future Air Force and DoD IO systems and weaponry. It is a shared, collaborative process that affords experts in IO and traditional warfighting disciplines an opportunity to discuss in an open-forum, working level environment their IO requirements, challenges, and solutions.

Phoenix Challenge's mission is straight forward. Phoenix Challenge is IO concepts and solutions exploration in action.

Phoenix Challenge was conceptualized and institutionalized by Mr. Thomas Sweet, a retired naval cryptologist and information warfare officer whose final active duty Navy assignment was at the AFIWC.

Phoenix Challenge allows the Air Force and other conference participants an opportunity to make a real difference in their organization and within the IO community. The information and technology shared at this conference is invaluable, and offers each participant a fresh perspective on the many challenges our Services and Department of Defense face on a day-to-day basis.

Mr. Thomas Sweet

The most recent Phoenix Challenge conference, scheduled April 22–25, 2002, was the fifth in a continuing conference series, and the first to be aligned with the Armed Forces Communications and Electronics Association (AFCEA) Fiesta Technet 2002 Conference that was held in San Antonio, Texas. Typical of its previous conference agendas, there was a day set aside for DoD and Federal government to share their IO perspectives on warfighting and homeland defense. The conference identified and addressed a myriad of IO challenges and solutions that transcend both service and organization boundaries, and offered industry and academia an opportunity to bring forward their innovative, state-of-the-art IO technologies and capabilities to a large DoD and Government audience. In an effort to increase operational and technical relevancy to the Phoenix Challenge program and conference series, AFIWC added a "warfighter day" to allow participants an opportunity to gain a better appreciation and understanding of the IO challenges and problem sets facing Unified and Specified Commanders in the performance of their assigned operational and warfighting missions.

In a recent disclosure and in large part due to Phoenix Challenge, the Office of the Secretary of Defense for Command, Control, Communications and Intelligence (OSD/C3I) has programmed substantive funds in the Fiscal Year 2003 Five Year Defense Plan, or FYDP, to develop a DoD-wide IO technology and requirements database to ensure information about IO technologies and capabilities are available on demand and in near real time to DoD and Government agencies involved in IO system and weapons development. The AFIWC will serve as the Executive Agent for this project.

We have come a long way since our first conference back in November of 2000, and the reason for this success is directly tied to the over 275 participants from 130 DoD, Government, Industry, and Academic organizations who routinely participate in Phoenix Challenge conferences. These are folks who share common goals and interests, and who really want to make a difference in how our services, our Government, and our nation conducts Information Operations in the 21st Century.

Mr. Thomas Sweet

For more information on Phoenix Challenge and how you can participate in upcoming conferences, you can visit the Phoenix Challenge public Web site at http://afiwcweb.lackland.af.mil or by contacting the primary Phoenix Challenge points of contact at sandra.vasile@lackland.af.mil or thomasa.sweet@lackland.af.mil.

We are identifying concepts and leveraging potential technology solutions to some of the more critical challenges and problem sets faced by the Air Force and DoD in an effort to execute a successful IO campaign against a more technologically advanced and increasingly more asymmetric threat.

Ms. Sandra Vasile
Program Sponsor

Software Decoys for Software Counterintelligence

by Dr. Neil C. Rowe, Dr. J. Bret Michael, Dr. Mikhail Auguston, and Mr. Richard Riehle

Some information systems are critical to defend against malicious attack. Yet they often rely on just the same countermeasures as any system—firewalls, authentication, intrusion detection systems, and encryption—although politically motivated attackers may be far more determined than hackers to bring them down. Future information security will increasingly use ideas from military defensive tactics[3] to effectively defend critical information systems. This will include automatic "counterintelligence" with deliberately deceptive behavior, what we call "software decoys." Decoys can deceive attackers into thinking their attacks have succeeded while protecting key assets at least temporarily.

Much good work has been done on intrusion detection systems,[8, 11] but only recently has there been corresponding work on how an attacked system should respond. Many respond to serious attacks by turning off the network connection, a high cost in today's networked world. Such a response tells the attacker they have been detected and this may just direct them to better targets. Moreover, defenders lose information about how the attack would have proceeded which they could have used to make defense more effective.

A cyber-attack is an attack on resources to gain tactical or strategic advantage just as in regular warfare. Deceptive responses can be automated much like the attacks. Such responses can be very effective because attackers depend on the honesty of the computer systems they attack. Deception can confuse their planning or frustrate them for a while without giving away our recognition that we are being attacked. This could be especially important during intensive information warfare when terrorists attempt to bring down critical systems in a short period of time: Delay permits time to analyze the attack and plan a response. Deception also allows us to turn an attacker's own strengths of patience and determination against them, much as Asian martial arts like Akido do with physical attacks.

The Concept of a Software Decoy

We have been researching "intelligent software decoys"[9]. We use this term to cover a spectrum of deceptive defensive activity[1]. This can range from mimicking normal behavior of the computer system (as when an attacker thinks they have gained system administrator privileges and we pretend they can modify key directories), through inventing appealing activities for the attacker (as when an attacker overflows a buffer and we pretend they have changed the behavior of the operating system), to new facilities (as when an attacker gets clues to a trap site with apparently vulnerable software). Appropriate deceptive tactics depend on the value of the resources being protected and the danger of the attack. But the general idea is to limit or confine[7] attacks that get through our first line of defense rather than stop attacks. Decoys differ from honeypots[4] in providing defense, not data.

Decoys are easiest to make when simple effects (like denial-of-service) are sought by attackers. They will generally work best against hands-on adversaries as opposed to automated scripts, though unpredictable responses by a decoy could well foil a script. Effective decoys need not be complex. Simple ploys in warfare can be surprisingly effective when their timing is right, they are consistent with enemy expectations, and they have some creativity.

Decoying capabilities should be distributed through an operating system and applications programs to provide a uniform front to attackers with no single point of compromise. They could go in Web servers, mail servers, and file-transfer utilities to address denial-of-service attacks and attempts to jump into the operating system. They could go in directory-listing capabilities to provide false information about sensitive directories. They could also go in network routers to address denial of service and suspicious patterns (like strings of nulls) with connection errors. More ambitious decoys could be embedded in all file-writing capabilities or all security-related activities of the operating system, through the use of "wrapper" technology that automatically inserts checking code around sensitive statements ("instruments" it).[9] While this may sound ambitious, an analogous technology exists for instrumenting code to calculate software metrics and monitor software at runtime, and such instrumentation has been successfully accomplished for large software systems—it is not hard for simple open-source operating systems like those for small devices.

We can distinguish levels of decoying. At the simplest level are memoryless decoys that respond the same way to the same local context. A behavior model based on an "event grammar" can operate on the system log to detect suspicious local context. It can use sophisticated ideas from the field of temporal logic. Creativity of decoy responses can be done with generative grammars having random choices. For instance, we have written generators for fake error messages (like "Error at 2849271: Segmentation fault") and for fake directory listings (with fake file names, dates, sizes, and subdirectories).

At one level, a decoy can remember other invocations of the same code. For instance, a server can store details of other transactions it has serviced so that it can recognize denial-of-service attacks. At another level, decoys in different software modules can share information about an attack, as when an attacker installs their own operating system. Finally at the highest level, a decoy can simulate the entire operating system itself within a "sandbox" or safe environment. This would be helpful when Trojan horses of unknown capabilities have been inserted into an operating system and the decoy must simulate them.[2] The higher levels of decoying require an architecture of response management.[6]

Types of Software Decoy Responses

A generally useful decoy tactic is the exaggeration of intended attacker effects: *Good deceptions should confirm preconceptions of the deceived*. Under a denial-of-service attack, for instance, we can pretend to increase the load on a computer system by deliberately delaying system responses. This can be done by calculating and implementing delays, accomplished by additional process-suspension time, into the servicing of attacker transactions,[10] with perhaps additional scripted interaction with the attacker.

An important factor in this is the probability that we are under attack. Unfortunately, new vulnerabilities and new techniques for exploiting those vulnerabilities are constantly being discovered. Recent hacker behavior shows an increasing automation of attacks, increasing use of rootkits, decreasing use of probes, and an increasing use of encryption for network communication.[4] But a determined adversary like a terrorist group will want to try new methods we have not anticipated. We must then use general principles to estimate the probability we are under attack, and respond proportionately to this probability. We can use current intrusion detection methods for this, both anomaly and misuse detection, but an especially helpful clue we are investigating are reports from similar sites about attacks that they are undergoing.[5] Automatic data mining from system logs can be helpful at those sites to analyze how they were attacked.

Decoy delays can be accomplished by process-suspension time alone, but alternatives can make the deception more interesting and engaging to the attacker. Many attackers see their activities as like playing a computer game, so some game-like behavior in the decoy could be helpful, as could "showmanship."[1] This could involve user interactions such as requests for authorization, requests to confirm allocation of more system resources like memory, deliberate errors, and invocation of new scripts pretending to be system-administrator tools—we can get creative.

Responses of software decoys must necessarily vary with resources available to fight the attack. Consider denial-of-service decoys for transaction servers. Delay exaggeration is only effective below a certain system load, because good deception requires that we still process the transactions albeit more slowly. If the attack intensity continues to increase, we could systematically simplify transactions, without telling users, by ignoring less important parts of the input. Or we could respond to a transaction with a cached result of a similar transaction, an effective idea for denial-of-service attacks doing the same transaction repeatedly.

If the attack intensity continues to grow, the system has no choice but to refuse transactions. However, we may still fool an attacker if we substitute a low-resource interaction that could conceivably result from a successful attack. For instance, we could say "Buffer overflow" and start what appears to be a debugger with "Stopped at line 368802 of module serv89—singlestep?" Or we could claim memory needs to be reallocated due to the high system load, and give the attacker a fake opportunity to change module memory requirements. Eventually however, if attack intensity continues to increase we must turn off the network connection and terminate the game with the attacker.

Attackers will eventually recognize decoys, and will plot countermeasures such as ignoring sites with recognizable decoy "signatures." But we can plot to counter the countermeasures, and so on. The classic field of game theory provides methods to analyze such situations and find our best overall strategy.

Endnotes

1. Bell, J. B., & Whaley, B., *Cheating and Deception*, New Brunswick, NJ, St. Martin's Press, 1991.
2. Bressoud, T. & Schneider, F. B., "Hypervisor-based Fault-tolerance." *ACM Transactions on Computer Systems*, Vol. 14, No. 1, pp. 80–107, Feb. 1996.
3. Fowler, C. A., & Nesbit, R. F., "Tactical deception in air-land warfare," *Journal of Electronic Defense*, Vol. 18, No. 6, pp. 37–44 & 76–79, June 1995).
4. The Honeynet Project, *Know Your Enemy*, Boston, Addison-Wesley, 2002.
5. Ingram, D., Kremer, H., & Rowe, N., "Distributed Intrusion Detection for Computer Systems Using Communicating Agents," *Proceedings from the 6th International Symposium on Research and Technology on Command and Control*, Annapolis, MD, June 2001.
6. Lewandowski, S., Van Hook, D., O'Leary, G., Haines, J., & Rossey, L., SARA: "Survivable Autonomic Response Architecture," *Proceedings from DARPA Information Survivability Conference*, Anaheim CA, June 2001, Vol. 1, pp. 77–88.
7. Liu, P. & Jajodia, S., "Multi-phase Damage Confinement in Database Systems for Intrusion Tolerance," *Proceedings from the 14th Computer Security Foundations Workshop*, Cape Breton, NS, pp. 191–205, June 2001.
8. Lunt, T. F., "A Survey of Intrusion Detection Techniques," *Computer and Security*, Vol. 12, No. 4, pp. 405–418, June 1993.
9. Michael, B., Auguston, M., Rowe, N., & Riehle, R., "Software Decoys: Intrusion Detection and Countermeasures," *Proceedings from the 2002 Workshop on Information Assurance*, West Point, NY, June 2002.
10. Somayaji, A., & Forrest, S., "Automated Response Using System-call Delays," *Proceedings from the 9th USENIX Security Symposium*, pp. 185–197, August 2000.
11. Vigna, G. & Kemmerer, R. A., "NetSTAT: A Network-based Intrusion Detection Approach," *Proceedings from the 14th Annual Computer Security Applications Conference*, Scottsdale, AZ, pp. 25–34, December 1998.

Biographies

Dr. Neil C. Rowe is Professor and Associate Chair of Computer Science at the U.S. Naval Postgraduate School where he has been since 1983. He has a Ph.D. in Computer Science from Stanford University (1983), and E.E. (1978), S.M. (1978), and S.B. (1975) degrees from the Massachusetts Institute of Technology. He has done research on intelligent access to multimedia databases, information security, image processing, robotic path planning, and intelligent tutoring systems. He has authored over one hundred technical publications and a book. He may be reached at ncrowe@nps.navy.mil.

Dr. James Bret Michael has been Associate Professor of Computer Science at the U.S. Naval Postgraduate School, Monterey California since 1998. He received his M.S. (1987) and Ph.D. (1993) degrees from the School of Information Technology and Engineering at George Mason University, and B.S. (1983) from West Virginia University. His research interests include both information operations and computer security for distributed computing systems such as those in missile defense, has authored over fifty technical publications and a book, and is a senior member of the IEEE. He may be reached at bmichael@nps.navy.mil.

Dr. Mikhail Auguston is an Associate Professor of Computer Science at New Mexico State University. He has graduated summa cum laude in Mathematics from University of Latvia in 1971 and received Ph.D. in Computer Science from the Glushkov Institute of Cybernetics in Kiev (USSR) in 1983. He has more than thirty years of research experience in programming language design and implementation, program testing, and debugging tool design, and has authored more than sixty technical publications. He may be reached at mikau@cs.nmsu.edu.

Mr. Richard Riehle is a Visiting Professor of Computer Science at U.S. Naval Postgraduate School. He is also a Principal of AdaWorks Software Engineering, a consulting firm that specializes in software development and training in Ada. Mr. Riehle has over twenty-five years in software development in both military and non-military systems. His current interests are software architecture, software deception techniques, programming language design, and software reliability. He has a B.S. from Brigham Young University and an M.S. in Software Engineering from National University. He may be reached at rdriehle@nps.navy.mil.

We must keep the fire lit…
the warfighters depend on us!

Information Assurance Campaign: Keeping the Fire Lit in 2002

Last year we put a blitz on Information Assurance awareness and eliminated a multitude of vulnerabilities through the Information Assurance campaign. Twelve monthly themes focused our attention on a succession of important IA issues: from *Roles and Responsibilities* to *Threats and Countermeasures*…from *Digital Devices* to *Computer Network Defense*…from *Web Security* to *Information Assurance in the Expeditionary Aerospace Force.* Our collective knowledge in these areas has significantly improved. Many network vulnerabilities were also eliminated through aggressive problem identification and resolution, and the use of Information Assurance tools.

We covered much ground last year, but our campaign is far from complete. More awareness and network protection actions are necessary before we can declare victory for the Information Assurance campaign.

In 2002, we must focus on Information Assurance activities directly supporting the war on terrorism. Our new campaign theme reflects this focus: *Defeating Global Terror…Demands Effective Information Assurance.* This year, several operationally-oriented Information Assurance themes have been planned, including *Contingency Planning, Operational Security, Remanence Security (Sanitization and Destruction),* and *Vulnerabilities and Incidents.* We will also revisit some important 2001 themes, such as *Web Security, User Responsibilities,* and *E-mail.*

I highly encourage everyone to become fully engaged in the continuing Information Assurance campaign. We must keep the fire lit…the warfighters depend on us!

Lt Gen John L. Woodward Jr.
Air Force Deputy Chief of Staff for Communications
and Information, Washington, D.C.

Information Systems Se

by Mr. Gordon Steele, IATAC Deputy Director

Many companies today have spent time and money on their Internet sites by investing in defenses against computer security incidents. Despite the best planning, incidents do happen and defenses are overrun. When that occurs an incident response capability may be all that stands between an enterprise's computing environment and an incident that can threaten even the viability of the enterprise.

Information Systems

There has been an increasing trend in recent years to extend enterprise-computing environments to the desktop by deploying distributed computing solutions. *Distributed computing* is a relative term among industry professionals. Definitions range from the client/server computing model through fully distributed processing where data, data management, applications, end user interfaces, and end user devices all reside on different hosts that are linked by networks.

Widespread migration of the computer industry toward distributed computing has had a significant impact on information system security. Multiple information systems may criss-cross a single host, making compromise of that central host a potentially lucrative target for attackers. "Single sign-on trust" domains—which enable Web surfers to visit multiple, unrelated secure servers after having entered a password just once on one site—also provide attackers the opportunity to compromise multiple hosts by finding the weak link in the domain. When an attacker finds a poorly patched or configured host, the attacker may use it as a springboard to attack other hosts that "trust" the now-compromised host.

A framework for depicting information systems in today's distributed computing environments is illustrated in Figure 1 (see page 15). In the framework the term "business process" means a related group of steps or activities that use people, information, and other resources to provide goods or services to internal or external customers. The framework consists of seven linked elements—

- Internal or external customers of the business process
- Products or services generated by the business process
- Steps in the business process
- Those who execute business processes (Participants)
- Information the business process uses or creates, and
- Technology the business process uses.[1]

This framework demonstrates that information systems touch many elements of an enterprise. Information systems affect, and are affected by each of their seven elements. Thus, any taxonomy for information system security incidents must consider all of the seven elements, not just the computers and networks that support them. This thinking is reflected in the reality that today's crime statistics report crimes committed by or committed against every element of information systems.

Some of the more traditional crimes, such as fraud, are often facilitated by information systems. When these crimes are discovered, an information security incident response team

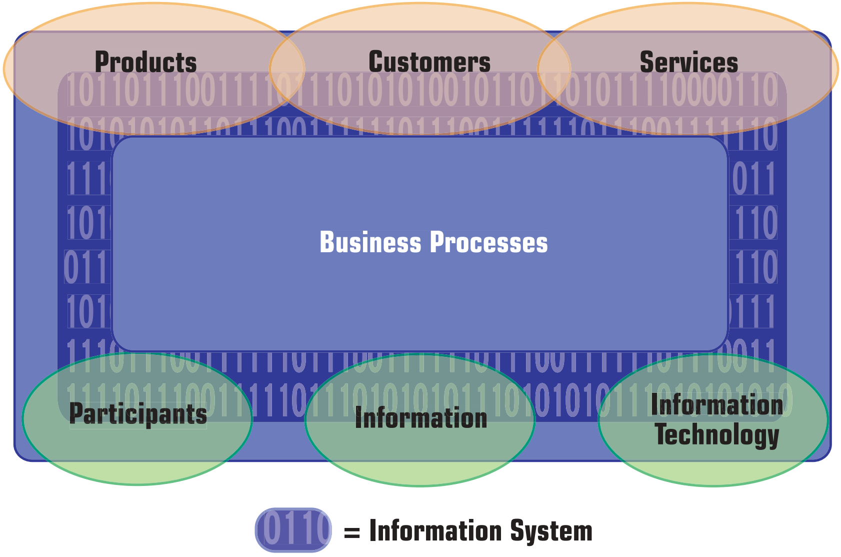

Figure 1. A framework for depicting information systems in today's distributed computing environments.

will most likely be involved in the resolution because an information system facilitated the crime.

Incidents and Incident Response

When people discuss computer security incidents they usually refer to one single aspect of the incident. They say, for example, that they were hacked, that their Web site was defaced, or any number of other things related to the way their site was attacked or the result of the attack. While these statements reflect certain aspects of multifaceted computer security incidents, they do not completely characterize them. Incident handlers require language that describes incidents more precisely and more fully. This is necessary to ensure that persons involved in incident handling address all aspects of an incident. Similarly, a common understanding of incident response methodologies is required by incident handlers to ensure that they are thorough and efficient in restoring the computing environment to a safe state. This article discusses both a common language for the incident handler, a widely practiced incident response methodology, and puts these together to describe a notional incident/incident response flow.

Incident Taxonomy

The IEEE Standard Dictionary of Electrical and Electronics Terms defines taxonomy as a classification scheme that partitions a body of knowledge and defines the relationship of the pieces. Multiple incident taxonomies are in the literature, but to date the most satisfying is that which was proposed by Howard and Longstaff in their Sandia National Laboratories report entitled, *A Common Language for Computer Security Incidents* (http://www.cert.org/research/taxonomy_988667.pdf). The taxonomy is depicted in Figure 2 (page 17). Read from left to right it could be interpreted—

an attacker uses a tool to exploit a vulnerability by performing an action against a target that results in an unauthorized result for the pur-

pose of achieving one or more objectives.

The two items in green, "action" and "target" together make up an event. An event can be thought of as any significant occurrence in an information system that requires users to be notified or an entry to be added to a log. Actions are steps taken by a user or process in order to achieve a result. Actions are directed against logical or physical entities known as targets. Examples of actions include probing, scanning, reading, modifying, etc. Examples of logical targets include accounts, processes, and data. Examples of physical targets include components, computers, and networks.

Taken out of context events are neither good nor bad. There are times when specific events might be authorized, such as when an administrator conducts a vulnerability scan of his or her network to identify exposures that should be reduced. However, when a tool that exploits a vulnerability performs the action, that event becomes part of an attack. Attacks are concerted efforts by an attacker to achieve an unauthorized result. The term "tool" is used here in a broad sense. It is a means of exploiting a vulnerability. Thus tools include anything from a set of commands input by a user to autonomous agents like worms. Vulnerabilities are flaws in the design, implementation, or configuration of hardware and software that may cause an information system to behave in an unpredictable, insecure manner. Attackers attempt to achieve unauthorized results such as increased access, disclosure of information, corruption of information, denial of service, and theft of resources by using tools to exploit these vulnerabilities.

Attacks do not occur without a reason, and they do not occur in isolation. The reality of Internet attacks is that attackers often do so repeatedly using multiple methods until they have achieved some objective. For example, popular hacking doctrines (e.g., those taught by the SANS Institute and by Foundstone) prescribe the following steps in an attack—

- Conduct reconnaissance
- Scan
- Exploit systems
- Gain access
- Elevate access
- Conduct application-level attacks
- Launch denial-of-service attacks
- Keep access
- Cover your tracks

In the context of the attack taxonomy above these "attacks" can be seen as elements of a methodical campaign involving multiple attacks by multiple means to achieve some objective. Thus, Howard and Longstaff define an incident as—

A group of attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing.

This is why it is both inaccurate and incomplete to characterize an incident in terms of just a particular attack technique observed, or a particular result achieved. It is conceivable, even likely, that if one attack is observed it is only part of a concert of attacks that are underway to achieve an objective beyond the immediate result achieved by the present observed attack. Thus the incident handler must deal with all elements and aspects of an incident, from attacker through objective.

Incident Response Methodology

Just as attackers follow a fairly defined methodology, incident handlers do as well. Perhaps the most popular incident response methodology today is taught by the SANS Institute. It consists of the following six phases—

- Prepare
- Identify
- Contain
- Eradicate
- Recover
- Follow-Up

This notional construct is suitable for guiding the workflow for most incident response operations, however it should be tailored to meet the technical needs and concerns of specific scenarios. This methodology is documented in the following exhibit. Each of the phases is described in the following section and corresponds to Figure 3 (page 18).

Phase 1: Prepare

In this phase of incident response an organization prepares for handling incidents. How well an organization has prepared in advance for computer security incidents will greatly affect how well its response personnel understand the extent of the incident; protect their information systems and in particular, sensitive data; and support investigations intended to result in an administrative, civil, or criminal remedy. Organizations that

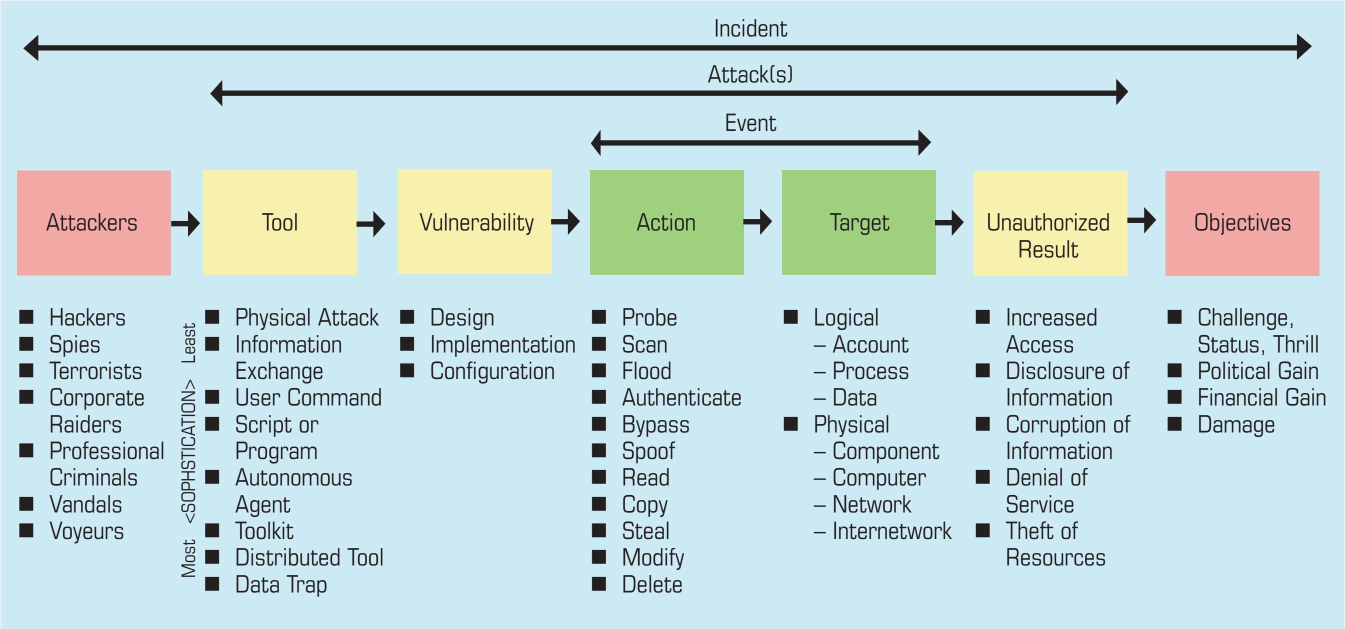

Figure 2. An attacker uses a tool to exploit a vulnerability by performing an action against a target that results in an unauthorized result for the purpose of achieving one or more objectives.

formalize their incident response capabilities often vest them in an incident response team. These teams go by various names in different enterprises, ranging from Computer Emergency Response Team (CERT), as they are referred to within the Department of Defense, to Computer Security Incident Response Teams (CSIRT), as they are referred to in many corporations.

During this phase the CSIRT should develop management support for an incident handling capability as well as select incident handling team members that will compose the core of the team. Team members should be trained in incident handling. An emergency communications plan should be developed that does not rely on the computing environment the team is tasked to protect. Other means of communication should be relied upon during an incident for two reasons. First, the network may not be available during an incident, and second, the network may not be trustworthy during an incident. Incident reporting facilities should be established and people selected to receive incident reports. Interfaces to law enforcement agencies and other CSIRTs should also be developed during this phase. Procedures for reporting incidents should be established for users.

The CSIRT should subscribe to a security alert service, or personnel should be designated to monitor security portals where emerging threats and vulnerabilities are being discussed in real time for the platforms in the organization's computing environment. Useful portals include the SANS Institute's Internet Storm Center, the CERT Coordination Center (CERT/CC), the four principal antivirus vendor Web sites and Bugtraq. Subscription services come in varying degrees of usefulness from free E-mail alerts to 24x7, "call-you-at-home," advisory services that are available from some managed security services companies. Membership in the Forum of Incident Response and Security Teams (FIRST) carries with it the extremely helpful benefit of chatting securely with incident handlers in other teams as Internet incidents unfold. As with anything, the quality of the product is often directly related to your investment. When it is perceived that there is potential risk to the organization from an emerging threat or recently discovered vulnerability, an overall assessment of risk to the organization should be conducted. The risk assessment relates the likelihood that a security incident related to the operational threat will occur (if it has not already), to the damage that will result if it does. Risk to the organization attributed to a threat should be categorized as high, medium, or low.

Criteria for classifying the severity of incidents should be developed during this phase. Unlike the criteria developed

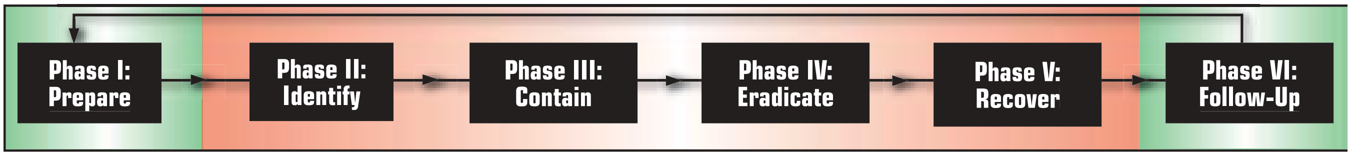

- Conduct mission preparation
- Establish and post warning banner
- Develop management support for an incident handling capability
- Select incident handling capability
- Develop an emergency communications plan
- Provide easy reporting facilities
- Conduct training for team members
- Establish guidelines for inter-departmental cooperation
- Pay particular attention to relationships with system administrators and network managers
- Develop interfaces to law enforcement agencies and other Computer Security Incident Response Teams

- Receive incident notification
- Assign a person to be responsible for the incident
- Determine whether or not an event is actually an incident
- Be careful to maintain a proper chain of custody
- Coordinate with the people who provide your network services
- Notify appropriate officials

- Deploy the on-site team to survey the situation
- Keep a low profile
- Avoid if possible potentially compromised code
- Backup the system
- Determine the risk of continuing operations
- Continue to consult with system owners
- Change passwords

- Determine cause and symptoms of the incident
- Improve defenses
- Perform vulnerability analysis
- Remove the cause of the incident
- Locate the most recent backup

- Restore the system
- Validate the system
- Decide when to restore operations
- Monitor the system

- Develop a follow-up report

Figure 3. Incident Response Methodology Phases

above for assessing risk from a potential threat, these criteria are applied when the threat actually manifests itself. These severity criteria are key to determining the level of effort the organization applies to the response effort.

One additional consideration related to incident severity that is often overlooked is that the team should have a definition of when to declare a security disaster and thereby invoke the organization's continuity of operations plans (COOP). The criteria for declaring a disaster may vary depending upon the organization's reliance on the affected information systems.

In addition to the operational, management, and administrative controls described thus far, certain technical controls should be put in place during the preparatory phase. For example, a central timeserver should be deployed in the environment because it can be particularly difficult to correlate events across a network when each computer maintains its own system time. A common time reference simplifies this considerably. If DHCP services are deployed, the capability to maintain a lasting record of IP address leases should be developed. Normally records of IP addresses are transient and may not be available to support an incident response or an investigation for long unless provisions have been developed to do so. One or more central log servers should be established on a fortified bastion host. Ideally, the enterprise antivirus server centralized logs would be stored here as well. In the event a worm with a file destructor payload propagates through the organization, these may be the only logs that survive. Connections to these servers should be limited to specific reporting hosts and all connections to them should be secure. Data stored on these servers should include event logs, security logs, application logs, and any other perishable logs that might be useful during an investigation. It is desirable that arriving log data be not only saved but spooled to a line printer as well. Anti-virus software should be installed and running on all of the organization's desktop computers.

Often CSIRTs will be called upon to support an operational recovery. Occasionally however, a CSIRT is called upon to respond in a way that preserves the organization's ability to pursue an administrative, civil, or criminal remedy. In such cases, the team will need to be prepared to acquire, analyze, and store evidence from computers in a forensically sound manner. Much of that evidence is perishable, and it is probably distributed throughout the network, if not throughout the Internet. Thus, when dealing with networked computers, one should not rely on just one evidentiary source such as the victim host. A sophisticated perpetrator will alter or delete logs on it. However, he or she may have left an audit trail in various places around the network. He or she may not be able to alter or delete all of them because he or she: is not aware of all of them; is not able to access all of them; or does not have time to access them. Therefore, incident handlers may find discrepancies as they examine the network. These discrepancies can be used to discern what actually occurred. Thus, given today's distributed computing environment, the evidence of an incident can be distributed throughout a computing environment. To prepare for these situations, the concept of evidence maps has been developed. These are lists that indicate where evidence is likely to be in the event of an incident. They are prepared in advance of an incident in order to assure expeditious acquisition of evidence. Thus, there are many sources of evidence to support the investigation—so many in fact that another plan should be generated in advance detailing how to **correlate** the data from these disparate sources.

Finally, incident handlers should know well the environment they will operate in. They should be very familiar with the major applications and general support systems that are deployed, and they should be able to map them to the network's logical and physical topologies. They should also know what constitutes normal activity so as to be better equipped to identify anomalous activity.

Phase 2: Identify

The indication that an event has occurred may come from myriad sources. Events may be reported by firewalls, IDS', file integrity checkers, etcetera. Regardless of how the alert is generated, information about the event must be analyzed so that a determination can be made as to whether the event is a computer security incident. Not surprisingly, some events initially reported as computer security incidents are determined to be technical performance issues on further examination.

Once it is determined that the event is indeed a computer security incident, the event must be subjected to triage so that the risk to the organization and the severity of the event is determined. The organization should use the incident severity categorization guidance developed in Phase 1 to accomplish this. If the event is severe enough to warrant invocation of the organization's computer security incident response capability, a team is then activated, and a team leader designated. The interdisciplinary nature of incident response should be considered at this point. The team should embody the technical skills required to contain the attacker(s) and to eradicate the exploited vulnerabilities, the attack vector, and any residual attack tools. But the team may also have to deal with diverse tasks requiring expertise in legal matters and public relations, particularly if the incident has spread from this organization to other sites.

A determination should be made at this point whether the organization wishes to proceed in a manner that preserves its ability to pursue an administrative, civil, or criminal remedy, or just pursue restoration of a secure operational state. If the answer to the former is yes, then the team must approach the computing environment as a crime scene. If pursuit of criminal action is envisioned this may be an appropriate point to notify law enforcement authorities and appropriate regulatory authorities. Note that evidence maps, if prepared in Phase 1, can be very helpful in obtaining perishable evidence in an expeditious manner.

Phase 3: Contain

At this point, the CSIRT is deployed and/or establishes a remote connection to the affected site to survey the situation. If deployed, the team will most likely deploy with a suite of tools on flyaway platforms. If affected systems have not been backed up, now is the time to consider doing so, however it should be noted that such backups should be quarantined as they may contain malicious code. The CSIRT should determine the risk of continuing operations at this point, and make a determination whether severing any host's connectivity to other hosts, up to and including isolating the organization from the Internet community.

It is very important for the team to build awareness of the threat throughout this phase. This includes an awareness of the technical details of the threat as well as an awareness of what the threat is doing within the organization's computing environment. External awareness is achieved by monitoring antivirus vendor sites, the NIPC, CERT, and SANS Web sites, etc. as well as coordinating with other teams such as members of the Forum of Incident Response and Security Teams (FIRST). Internal awareness is a considerably more complex undertaking. Event data may be generated by many internal sources including firewalls, intrusion detection systems, anti-virus engines, etc. The technical data describing exactly what is happening in the environment is a both a boon and a bane in that the amount of data that can be produced by all these event generators can be astounding. The team must therefore, come fully equipped to perform log collection, normalization, and analysis using automated tools. From this data, the CSIRT must determine which events are incident related, and which events can be correlated with other events.

Frequently, the team will perform a mini-vulnerability assessment of the affected site during this phase of the incident to quickly determine what vulnerabilities are exploitable, and then take precautionary steps using available network controls to reduce vulnerabilities found. Techniques for assessment may include the use of vulnerability scanners, port scanners, and war dialers. One benefit of the trend toward defense in depth is that it affords the CSIRT the ability to respond in depth. Current technical security controls may provide numerous choke points in the environment at which the team may exercise control over the incident. How they may be used to contain the incident is dependent upon the attack vector, so maintaining a growing awareness of the attacker's capabilities is paramount during an incident. Typical controls are described below—

- **Router**—Limit access to and from the site (and/or enclaves within the site) by IP address or port through the use of access control lists (ACL). On some routers, traffic can be shaped which can limit the effects of a denial of service incident.
- **Application Proxy**—Allows filtering of content at the application layer.
- **Firewall**—Provides the ability to drop or reject inbound and outbound packets based upon IP address, port, etc. May also be used to launch script in response to certain scenerios.
- **Remote Access Server (Dial-In and VPN)**—Limit connections to those possessing the credentials required for authentication. During an incident all remote access accounts should be reviewed.
- **Mail Transfer Agent (MTA)**—May allow filtering the content of inbound and outbound messages, including disallowing of file attachment types based on file extension.
- **AntiVirus Engine**—Allows scanning of data on media and in memory for signature matches with known malware, and heuristic scanning to identify suspicious behavior.
- **PC Operating System**—Allows limiting and/or disabling of nonessential services, file shares, incoming connections by port, etc.

The CSIRT will assess the situation and plan a course of action appropriate to the situation. Many options are available for response. These include restoring operations, online response versus offline response, involving public relations, identifying the attacker, prosecuting the attacker, as well as many others. The response strategy chosen will determine what actions will be taken, and consequently, what types of resolution are possible. The type of attack and the classification of the victim system will be considered in determining a response strategy. For example, different strategies will affect the availability of the victim system differently. Therefore the number of persons relying on the system, the criticality of data on the system, and the effects of having the system offline for various lengths of time should all be considered. Whatever the options chosen, they should be in full consideration of as many aspects of the organization's mission needs as possible.

Phase 4: Eradicate

During this phase the team will make the final determination of the cause of the incident and take whatever steps are necessary to eradicate the cause from the environment. Using information derived from the vulnerability assessment conducted during the last phase the team may advocate and/or deploy additional technical controls, as well as administrative, operational, and management controls. If it was necessary to wipe systems during the last phase, attempts may be made to restore data to rebuild systems using whatever clean backups may be available.

Phase 5: Recover

During this phase, the team will take final steps to restore the environment to a secure state. If systems have been offline, the team will also make a final determination as to when it is appropriate to restore service. This includes restart of services that may have been terminated and reconnection of network connections that may have been physically severed. The team will monitor the state of the system for some time to look for anomalous activity that may signify that the incident cause has not yet been eradicated.

Phase 6: Follow-Up

During this last phase, the team will develop a final report that summarizes the operation and makes recommendations to preclude future recurrences. This may include recommendations to alter the security architecture, including security policies. The extent and format of such reports varies according to the needs of the organization and any external bodies that may be involved, such as insurance companies, law enforcement organizations, and regulatory oversight bodies.

Incident/Incident Repsonse Flow

Figure 4 (page 22) describes the flow of an incident and an accompanying response. It incorporates both the incident taxonomy and the incident response methodology described earlier in this article. Of course, this timeline is drawn at a high level of abstraction. In practice all these steps have to be applied using the administrative and technical controls that exist in the computing environment. The benefit, however, of documenting this flow at this level is to allow handlers to envision where they might be at any given point in time independent of the details of technology. This can help an incident handler to keep tabs on where they are in the incident process, where the attacker is probably going, and what their response options are at any given time. A timeline runs down the center of the exhibit. Time, *t* = zero is at the top of the exhibit. The incident taxonomy is depicted to the left of the timeline, and the corresponding incident response methodology is depicted to the right of the timeline. Each of the elements in the incident taxonomy and the incident response methodology has already been discussed. What is important to note however is that when the two are considered together, the additional opportunities to head off an incident become apparent. These opportunities are discussed here.

The first opportunity to prepare for an incident, beyond implementing the security architecture, is to learn about potential attackers before an incident occurs and tailor the enterprises' defenses to ward off threats from that attacker. For example, members of the security community occasionally lurk within IRC channels. While doing so they may observe hosts being compromised and thereby becoming zombies that are part of distributed denial of service (DDOS) networks. Lists of such zombies are traded between security teams through such organizations as FIRST. It behooves one to use the addresses from such lists to—

1. Try to contact the administrator for each host and have them eradicate the tools that have been installed on their machines, and
2. Limit your computing environments' exposure to inappropriate behavior from those hosts.

The next opportunity an incident handler may have to ward off an attack before it happens is when he or she becomes aware that an attack tool has been released that targets elements of information systems deployed within his or her computing environment. Specifics of the attack tool(s) should be researched at that point and the enterprise's risk posture should be reassessed. Then, appropriate controls should be put into place to counter any emerging threat from the tool(s). Next, when a previously unknown vulnerability is discovered it should be either corrected or compensated for immediately. There is a window of opportunity for attackers to exploit newly published vulnerabilities while administrators fix their networks. This is a cat-and-mouse game that can, and does, yield high rewards for attackers that are first to exploit newly published vulnerabilities. It is important to note at this point that this methodology applies to **all** elements of an information system. Thus, vulnerabilities can come from issues related not only to information technology, but from issues related to business processes, participants, information, products, services, and customers as well.

One last opportunity to counter an attacker comes while the attacker is actively using tools to detect or to exploit vulnerabilities, but before he or she has achieved his or her objectives. If detected soon enough, the attacker may still be in the early stages of a series of attacks that contribute to attainment of a larger objective. If thwarted at this point, a more serious incident may be avoided. Once that opportunity has passed however, the only opportunity to stop an incident from escalating in severity is an automated response that detects sus-

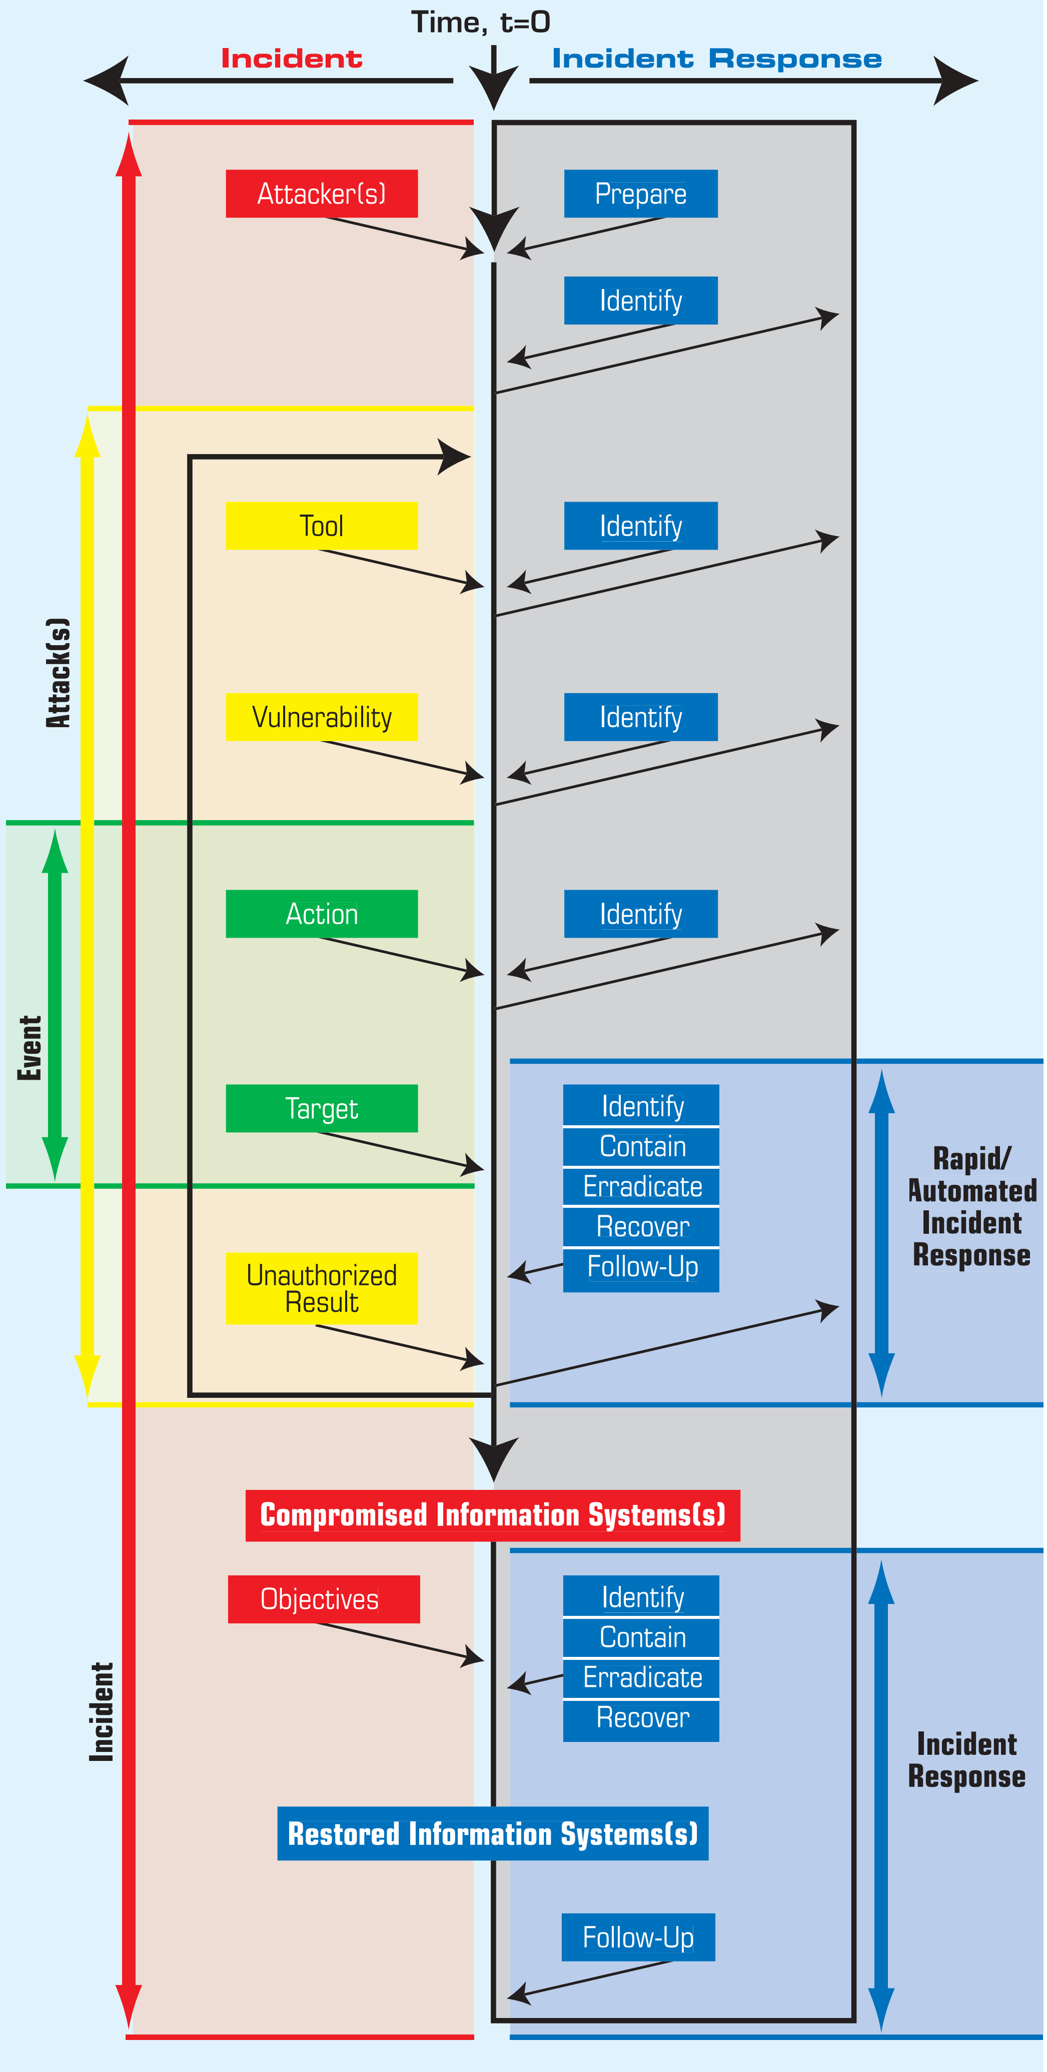

picious activity and employs automated response(s) such as blocking, session termination, account lockout, etc. From that point on however, the incident response team will probably be dealing with a compromised environment, and all of the steps required to contain, eradicate, and recover must be performed as described earlier.

Summary

This article has presented a taxonomy for incidents and a widely practiced incident response methodology. Future articles in this series will deal with establishing a computer security incident response capability at your site, as well as appropriate response measures for particular scenarios.

Endnotes

1. See: Steven Alter, *Information Systems, A Management Perspective* (2d ed.), Benjamin/Cummings Publishing Company, Inc., 1996.

Gordon Steele is currently the Deputy Director of the DoD Information Assurance Technology Analysis Center (IATAC). He may be reached at iatac@dtic.mil.

Figure 4. Incident responders need to be mindful of where they are in the incident taxonomy at any given point in time.

Call for IO Technology Exhibits

NIOW 02

Naval IO Symposium

Sponsored by the Fleet Information Warfare Center (FIWC)
June 11–12 • Little Creek Conference Center

NIOW 02 will be structured as an IO technology symposium and exposition. The objectives of the wargame are—

- Provide a professional forum for discussion, education and evaluation of current and future Joint and Naval IO issues

- Provide IO professionals with an opportunity to learn about the latest DoD and industry technology innovations and concepts in IO through interactive briefings, demonstrations and displays.

Demonstrations, briefs, and technology displays are being scheduled and the final agenda will be published in early June.

Intersted in Attending?

Seating will be limited to 120 attendees. Registration for the Symposium is available at http://www.fiwc.navy.mil.

Participation Fee: $10.00

Showcase Your Innovations & Concepts

This is a great opportunity to share your IO technology innovations and concepts. Interested parties should E-mail **iatac@dtic.mil** with papers/presentations of the technology/capability demo for consideration. Capabilities that are not necessarily solely IO, but could have some applicability to the IO mission are also welcome.

Technologies of Interest Include—

- Electronic Warfare (EW)
- Operational Security Decision Aides
- Effects-Based Operations
- Visualization Technologies for Predicting Attacks
- New Applications of IO Technologies
- Knowledge Management
- Mission Planning Tools

The symposium will be **GENSER SECRET**.

Exhibitors and attendees must hold a **SECRET** level clearance.

Deadline for Consideration: June 3, 2002

Continuity of Operations [COOP]

by Mr. Abraham Usher

On April 5, 2002, the National Information Assurance Partnership held a one day conference on Continuity of Operations at the headquarters of the National Institute of Standards and Technology.[1] The conference was co-sponsored by the Information Systems Audit and Control Association, the Association of Government Accountants, and PricewaterhouseCoopers, LLC. The purpose of this event was to "help promote the development of a more secure information technology infrastructure within the United States."[2] Mr. Newt Gingrich, former Speaker of the U.S. House of Representatives, presented a keynote address on the topic "Continuity of Operations: Planning and Response in Today's Environment." Three specific information assurance themes were common to all of the presentations and discussions during the day—

- Security must be built into information systems from conception to completion
- Modern information systems must be redundant and self-correcting
- Effective information assurance defenses require proactive, not re-active actions.

The notion that security can be "added on" to a system after it is complete is a fallacy. Many information system vulnerabilities are due to the fact that security was not a primary design consideration during system development. As part of his address, Newt Gingrich drew an analogy between building a house and building an information system. For people of most modern countries, indoor-plumbing is considered a common feature of homes and buildings. Not many people would have a house built from scratch, and after everything was completed, decide to include plumbing as an "add-on." Indoor plumbing is an important, integral component of building a house. Similarly, security architecture and design must be an integral part of building information systems, rather than an optional component that is added if money and time are left after the project is complete.

Disruptions to the availability of information system components must be anticipated and planned for as part of a comprehensive system architecture. Most experts agree that "if something can happen, it will happen" and that future attacks on our information infrastructure will come, it is just a question of when. To deal with such an all-encompassing, undefined threats, our communication and information systems must contain the following qualities—

- Self-adaptation
- Self-correction
- Self-healing
- Redundancy

An excellent high-level view of a system that exhibits such qualities is the Internet. The Internet is a loosely coupled network of networks that does not have a single point of failure. The routing protocols that allow traffic to flow from one autonomous network to another are adaptive and self-correcting in nature; if one node becomes non-functional, they can instruct their router to adopt a new pattern of traffic exchange. Similarly, if Internet routers that were unavailable become functional again, the Internet "heals" itself by recognizing these routers and including them in the pool of available nodes for packet exchange. A similar paradigm must be adopted with the information resources of commercial and government entities. Systems must not be limited to a single point of failure by any particular device or component, rather they must employ a distributed architecture capable of adapting to change. Redundant resources must be maintained both virtually, and geographically. In the case of September 11, off-site backups that were held near the Word Trade Center were rendered unavailable due to loss of power and communications. The assumptions regarding the area of physical and virtual effects of infrastructure attacks must be carefully weighed and considered. Leaders must create continuity of operations plans that address the secondary and tertiary effects of infrastructure attacks that result from external dependencies to other organizations and resources.

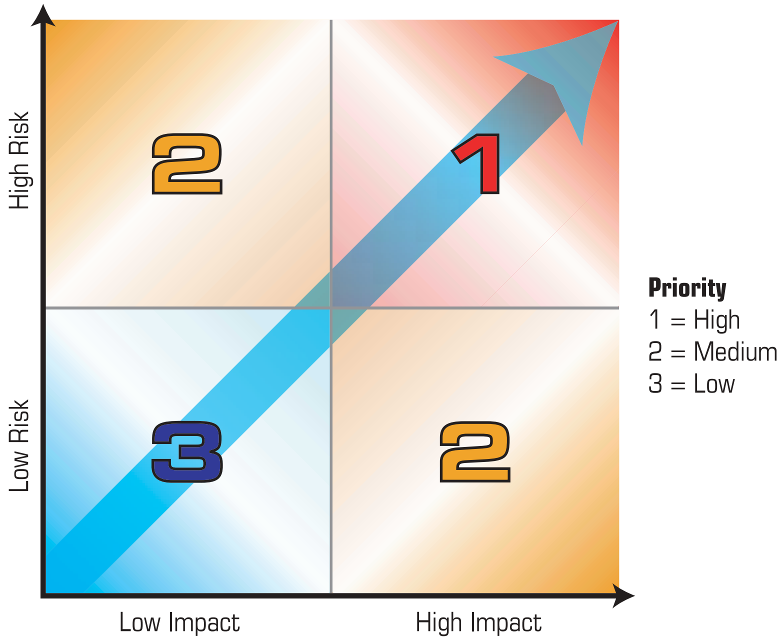

Figure 1. As risk increases, the danger or threat increases.

Chief Information Officers (CIOs), Chief Operating Officers (COOs), and other high level leaders face a difficult challenge in balancing the necessity of robust, redundant systems with the realities of limitations in capital and human resources. In order to provide the appropriate level of effort to the appropriate system or organizational capability, a process of prioritization is critical. A simple model for this type of prioritization is considering the relative impact of a system failure as well as the risk that the failure will occur (Figure 1).

Using such a model, decision makers and managers can focus their resources on system failures that are high-risk, high-impact in nature. System failures of a high-risk, low-impact or low-risk, high-impact nature could be categorized as the next level of importance in addressing and protecting. System failures of low-risk, low-impact nature may be documented, but might not require any actions or planning at all.

The nature of threats to the security and infrastructure of the United States are constantly evolving. Old military models of deterrence and passive defense are effective only when used against rational, nation-state entities. The concept of configuring defenses based on a known threat is futile when the risks faced are unknown. While operating in an environment of such uncertainty, the best way to defend critical resources and ensure continuity of service is to take a proactive approach to security. This entails creating systems that are created to effectively respond to specific threat capabilities, rather than responding to specific threats. For example in the case of computer viruses, software protection must be able to respond to the threats from current and future (unknown) viruses rather than just defending against known viruses.

More information on the National Information Assurance Partnership and the results of the Continuity of Operations conference are available by contacting the deputy director of NIAP, Terry Losonsky, at losonsky@nist.gov or 301.975.4060.

Endnotes

1. The National Information Assurance Partnership (NIAP) is a U.S. Government initiative designed to meet the security testing, evaluation, and assessment needs of both information technology (IT) producers and consumers. NIAP is a collaboration between the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA) in fulfilling their respective responsibilities under the Computer Security Act of 1987. The partnership, originated in 1997, combines the extensive security experience of both agencies to promote the development of technically sound security requirements for IT products and systems and appropriate metrics for evaluating those products and systems. The long-term goal of NIAP is to help increase the level of trust consumers have in their information systems and networks through the use of cost-effective security testing, evaluation, and assessment programs. NIAP continues to build important relationships with government agencies and industry in a variety of areas to help meet current and future IT security challenges affecting the nation's critical information infrastructure. More information is available at the NIAP Web site at: http://niap.nist.gov
2. As stated on: http://niap.nist.gov

Mr. Abe Usher is a research analyst at IATAC. He holds a B.S. degree from the United States Military Academy. He is currently pursuing an M.S. degree in Information Systems from George Mason University. Mr. Usher may be reached at iatac@dtic.mil.

What's New

Network Centric Warfare

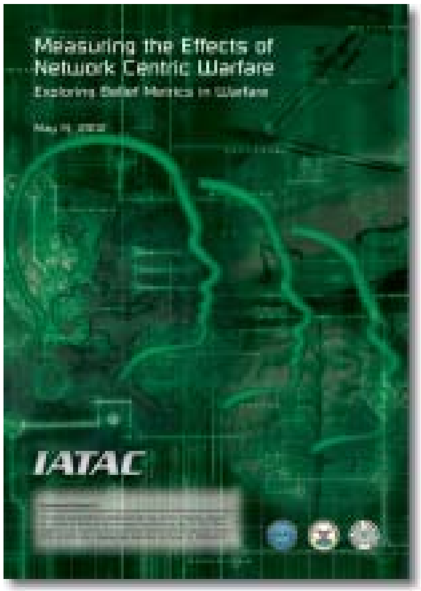

This report identifies key metrics for measuring the effects of network centric warfare (NCW). The goal is to present a series of quantifiable metrics that can be employed to measure NCW. The field of endeavor is a qualitative set of variables associated with the belief sphere of warfare. In the construct of this report, the belief sphere is divided into three distinct yet inter-related chapter headings: Unit Cohesion, Individual Morale, and Leadership.

In addition to detailed discussions, the potential effects of the metrics to both platform centric and NCW has been put in context in an applicable historical case study. Based on the development of these metrics, it appears that there are two phases in the implementation of NCW.

1. The Navy, and potentially the other Services, will build a comprehensive, linked network and will superimpose network centric capabilities onto its existing force architecture.
2. A new force structure will emerge, which will optimize these new concepts of warfare.

Key Attributes and Vulnerabilities of NCW

From a belief sphere perspective, NCW can affect many traditional variables that support the morale and cohesion of warfighters negatively. For example, the ability to operate in a more dispersed manner enabled by NCW runs counter to the wealth of evidence illustrating the importance of physical proximity to fellow soldiers/sailors/airmen, units, and leaders. Similarly, the notion of remote fire support or force protection and fighting with tailored, joint, ad-hoc units will challenge the unit cohesion and bonding that comes from time spent training and fighting together. The use of robotics and video teleconferencing will further challenge the warfighters' capacity to bond with their fellow Service men/women as well as with their commanders. Moreover, regardless of the particular metric or scenario, loss of connectivity would be disastrous for a network centric force. Though this would have a negative effect on all forms of Military operations, factors such as information loss would especially damage a network centric force that was trained and deployed to fight an NCW operation.

Findings

1. If the U.S. embarks upon building a network centric force, we must place the protection of the critical information flow at the top of the priority list.
2. The human aspect of conducting network centric operations will require new types of units, sailors/soldiers/airmen, organizations, and doctrines.
3. The key to fully developing the necessary personnel to realize an NCW operational capability is training. To create the new units and warfighters discussed above, we will have to train them to plan, exercise, and fight using new doctrine.
4. Operational concepts designed to protect platforms and equipment better may not be desirable from a human factors standpoint.
5. To understand the real benefits and vulnerabilities of NCW, the analytic community will have to re-evaluate the metrics used to determine effectiveness and which data to collect through experimentation and simulation.

Order Form

IMPORTANT NOTE: All IATAC products 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/regprocess.html.

Name ______________________________

Organization ______________________________

Address ______________________________

____

____

DTIC User Code ______________________________

Ofc. Symbol ______________________________

Phone ______________________________

E-mail ______________________________

Fax ______________________________

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

____

LIMITED DISTRIBUTION

IA Collection Acquisitions CD–ROM

❑ Fall 2001 ed.

IA Tools Report

❑ Firewalls (3rd ed.) ❑ Intrusion Detection (3rd ed.) ❑ Vulnerability Analysis (2nd ed.)

Critical Review and Technology Assessment (CR/TA) Reports

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

State-of-the-Art Reports (SOARs)

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

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

UNLIMITED DISTRIBUTION

IAnewsletters (Limited number of back issues available)

Volumes 1	❑ No. 1	❑ No. 2	❑ No. 3	
Volumes 2	❑ No. 1	❑ No. 2 (soft copy only)	❑ No. 3	❑ No. 4
Volumes 3	❑ No. 1	❑ No. 2	❑ No. 3 (soft copy only)	❑ No. 4 (soft copy only)
Volumes 4	❑ No. 1 (soft copy only)	❑ No. 2	❑ No. 3	❑ No. 4
Volume 5	❑ No. 1			

Fax completed form to IATAC at 703.289.5467

calendar

May 21-23
PACOM IA Conference
Honolulu, HI
Visit our booth
http://www.iaevents.com/Pacom/PacomNewInfo.html

30
Securing the Wireless Office (FNBDT Interoperability)
Information Assurance Technical Framework Forum
John Hopkins University, Applied Physics Laboratory, Laurel, MD
http://www.iatf.net/

Jun 3-6
3rd Annual DoD PKI Users Forum
San Diego, CA
http://www.iaevents.com/DoDPKI2002/PKINewInfo.html

11-12
FIWC IO Technology Symposium
Naval Amphibious Base, Little Creek, VA
SECRET/U.S. only
Call for IO Technologies—due 5/20/02, see ad on page 23,
E-mail: iatac@dtic.mil

11-13
TECHNET International
DC Convention Center
"Terrorism and Technology—The Critical Role of IT"
http://www.technet2002.org

12-13
Federal Information Superiority Conference (FISC) 2002
Sheraton Colorado Springs Hotel, Colorado Springs, CO
http://www.fbcinc.com/fisc/index.html

Jul 11
PKI Revisited—Information Assurance Technical Framework Forum
John Hopkins University, Applied Physics Laboratory, Laurel, MD
Limited to U.S. Government employees and U.S. citizens
http://www.iatf.net/

15-19
Information Warfare Seminar (IWS)
Secret (U.S. Only)
National Defense University
http://www.ndu.edu/irmc

17-19
3rd Annual IA Workshop
U.S. Military Academy, West Point, NY
http://www.itoc.usma.edu/workshop

27-Aug 2
SANSFIRE 2002
Boston, MA
IATAC will offer a presentation entitled "Do Borders Matter in Cyber Crime?"
http://www.sans.org/SANSFIRE02

IATAC

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