IAnewsletter

The Newsletter for Information Assurance Technology Professionals

Volume 8 Number 2 - Summer 2005

Common Technology Needs and Capability Gaps Across DoD's IA and CND Communities

also inside—

- IATAC Spotlight on Research—Dartmouth College
- The Kerf Toolkit for Intrusion Analysis
- IATAC Spotlight on Subject Matter Expert (SME)—Dr. Sergey Bratus
- Integrating Information Assurance into the DoD Acquisition System
- Threats Posed by and to 802.11 Wireless Networks
- Careless Keystsrokes Can Kill

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
Deputy Director:	Greg McClellan
Inquiry Services:	Peggy O'Connor

IAnewsletter Staff—

Creative Directors:	Christina P. McNemar Ahnie Jenkins
Art Director:	Bryn Farrar
Copy Editor:	Diane Ivone
Designers:	Kathy Everett Brad Whitford
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—
Winter 2005 October 3, 2005

Cover design:	Sergei Kalinin
Newsletter design:	Ahnie Jenkins

Distribution Statement A:

Approved for public release;
distribution is unlimited.

IATAC Chat

Gene Tyler, IATAC Director

This edition of the IAnewsletter introduces the Information Asssurance/Computer Network Defense (IA/CND) State-of-the-Art Report (SOAR) *and provides readers with background and insight into its creation.*

IATAC is pleased to announce our latest State-of-the-Art Report (SOAR): *Department of Defense (DoD) Information Assurance (IA) and Computer Network Defense (CND) Strategies—A Comprehensive Review of Common Needs and Capability Gaps*. As I write this Chat, the SOAR is undergoing its final technical review and will be published within days. This edition of the *IAnewsletter* introduces the SOAR and provides readers with background and insight into its creation. The article, Common Technology Needs and Capability Gaps Across DoD's IA and CND Communities, authored by Ms. Karen Goertzel, is a superb introduction to the SOAR. Karen's article explains the behind-the-scenes details that led to the creation of the SOAR, succinctly outlines the methodology used to prepare it, discusses the vetting process used to ensure that the SOAR can provide value to DoD IA seniors and planners, and highlights some next steps for the IA and CND communities. As the lead author of the SOAR, Karen is well suited to write the article. Because the content of the SOAR focuses on a sensitive topic and discusses planning capabilities, the SOAR will be released in limited distribution to authorized US government agencies and their contractors (Distribution C). You can access many of IATAC products on the Web (http://iac.dtic.mil/iatac/reports.html); however, reports such as the current SOAR can be accessed only by .mil or .gov addresses. I highly encourage all readers to view our Web site (http://iac.dtic.mil/iatac/) from time to time and to actively pursue a STINET account from DTIC http://www.dtic.mil/.

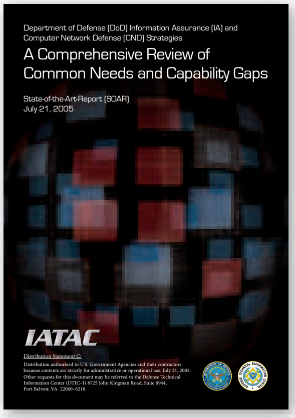
Department of Defense (DoD) Information Assurance (IA) and Computer Network Defense (CND) Strategies
A Comprehensive Review of Common Needs and Capability Gaps
State-of-the-Art-Report (SOAR)
July 21, 2005
IATAC
Distribution Statement C

IATAC SOAR:
A Comprehensive Review of Common Needs and Capability Gaps

The remainder of this edition includes additional information of value to IA professionals. It is no coincidence the SOAR article has a strong Research & Development (R&D) flavor. Please take the time to read our Chief Scientist's articles, IATAC Spotlight on Research—Dartmouth College and IATAC Spotlight on Subject Matter Expert (SME)—Dr. Sergey Bratus, Mr. Eustace King's Integrating Information Assurance into the DoD Acquisition System, and a question in the Letters to the Director column in which we address the Research & Engineering Portal. All of this information strongly applies to R&D communities and offers insight into the importance of IA and its relevance to all we do.

An added bonus in this issue of the *IAnewsletter* is a look at *Careless Keystrokes Can Kill*, the new video produced by the U.S. Strategic Command (USSSTRATCOM) and the Joint Task Force for Global Network Operations (JTF-GNO) which addresses the frightening fact that eighty percent of information needed to sabotage a mission can be obtained through unprotected, open sources. To inquire about copies of the video, please visit the Interagency OPSEC Support Staff—IOSS—at www.ioss.gov or the Joint Information Operations Center at www.jioc.smil.mil).■

Gene Tyler

Common Technology Needs and Capability Gaps Across DoD's IA and CND Communities

by Karen Mercedes Goertzel, CISSP

Across the US Department of Defense (DoD), a number of organizations have published strategies, plans, roadmaps, initiatives, and reference-capabilities documents, all in an effort to depict Defense-wide plans, requirements, and outstanding needs for Information Assurance (IA) technologies. These various documents can be said to generally fall into two areas: documents that characterize IA plans and requirements and documents that depict Computer Network Defense (CND) plans and requirements. Even though CND is formally acknowledged as a discipline within IA, as depicted in their strategic and planning documents, the focus and priorities of CND planners often differ significantly from those of broader IA planners. Moreover, even within the IA or CND discipline, there are often conflicts among the visions depicted in different organizations' strategic or planning documents.

This multiplicity of documents, all ostensibly containing complementary if not duplicative objects but reflecting different viewpoints, led the Information Assurance Technology Analysis Center (IATAC) Steering Committee to question whether it was possible to analyze the full range of DoD IA and CND plans and requirements contained in those documents to (1) reveal areas of unnecessary duplication and unexpected disjuncture and (2) to identify significant omissions. A team of IATAC IA Subject Matter Experts (SMEs) was tasked by the Steering Committee to perform an analysis of a broad, representative set of DoD IA and CND documents published by several different DoD organizations.

The findings of this analysis have been published as an IATAC State-of-the-Art-Report (SOAR). The main objective of the SOAR is to provide a synchronized portrayal of DoD, IA, and CND technology plans and requirements that highlights the gaps, overlaps, and omissions among them. This report is intended to provide those in DoD who are responsible for investing in technical research and development and those responsible for acting on the variety of sometimes conflicting plans with an informational basis that should help them more effectively establish their investment priorities and implementation decisions.

Objectives of the IA/CND SOAR

This IATAC SOAR Report is entitled Department of Defense (DoD) Information Assurance (IA) and Computer Network Defense (CND) Strategies—A Comprehensive Review of Common Needs and Capability Gaps. The report provides a summarized depiction of DoD's technical-capability requirements, challenges, and solution sets in Information Assurance (IA), including Computer Network Defense (CND). The information contained within the report has been gathered from a number of published and draft strategies, plans, roadmaps, initiatives, and reference-capabilities documents specified for analysis by the IATAC Steering Committee. Four main viewpoints, suggested by the Steering Committee, are represented in this SOAR: (see Table 1 on next page).

This SOAR necessarily reflects a "snapshot" of the source documents at the time at which they were reviewed and analyzed by the SOAR authors. This SOAR presents the needs and common requirements for IA and CND technical solutions captured in the source documents. The IA and CND Capability-Area Matrices are presented in a way that should enable the reader to easily characterize the individual technical solutions, solution sets, and Capability Areas according to the following criteria:

- Current status (deployed, planned to be deployed, or outstanding need)
- Specific technical capability provided by the solution or capability
- Individual group(s) of stakeholders invested in the solution or capability
- Range of stakeholders invested in the solution or capability

The SOAR is not intended to prioritize the outstanding needs identified by the various source documents nor does it provide extensive or in-depth recommendations or suggest comprehensive courses of action. The SOAR's content should, however, provide an informational basis for developing such recommendations and courses of action.

Research and Technology (R&T) Viewpoint	
1	This viewpoint focuses on longer-term research, development, and transition of IA and CND technologies. ■ Information Security (INFOSEC) Research Council (IRC) IA Hard Problems List (draft, Jan 2005) ■ Mitre Corporation (for NSA and the US Strategic Command) Assessment of IA/CND Focus Areas (Revised 1.0, Jun 2004) ■ NSA/STRATCOM CND Research and Technology (R&T) Program Manager's 1st Annual Report, Toward the Next Generation of Computer Network Defense, Version 4.0, Jun 2002 ■ NSA/STRATCOM CND R&T Program Manager's Second Annual Report, CND R&T Update and Program Manager's Plan, Draft Aug 2004
Information Assurance (IA) Component of the Global Information Grid (GIG) Architecture Viewpoint	
2	This viewpoint focuses on defining and planning for the short and medium term (GIG Increments 1 and 2) to implement and deploy IA capabilities required to secure the GIG. ■ National Security Agency (NSA) GIG IA Reference Capabilities Document, Volumes I and II, (Draft Version 1.0, 26 Oct 2004) ■ NSA GIG IA Capability and Technology Roadmap (Draft Version 1.0, 26 Oct 2004)
DoD IA Strategic Viewpoint	
3	This viewpoint focuses on broader strategic initiatives and issues associated with integrating IA capabilities into systems and operations (not necessarily specific to the GIG) throughout DoD, including issues related to capability definition, budgeting, and funding. ■ DoD IA Strategic Plan (Version 1.1, Jan 2004) ■ Joint Staff J6 Joint IA Campaign Plan (Final, 3 Dec 2004) ■ National Security Agency (NSA) Net-Centric IA Strategy (Draft Version 1.0, 30 Jun 2004)
DoD Computer Network Defense (CND) Planning Viewpoint	
4	This viewpoint focuses on short-, medium-, and long-term planning to implement and deploy throughout DoD technical and non-technical CND capabilities (as contrasted with broader IA). Planning issues include addressing short-term capability, development, and deployment schedules and budgets. ■ DoD CND Initiatives Program Plan (10 Jan 2005) ■ Office of the Assistant Secretary of Defense Network Information and Infrastructure) (ASD-NII) DoD CND Architecture Roadmap (Draft Dec 2003) ■ DoD Information Operations Roadmap (30 Oct 2003, Unclassified portions) ■ U.S. Strategic Command (STRATCOM) CND Strategy for Defense in Depth (28 Sept 2004) ■ Defense-wide Enterprise Security Solutions Group (ESSG) Finalized Working Group Meeting Minutes 14–16 September 2004

Table 1: Viewpoint documents for the IA/CND SOAR

Report development methodology

This SOAR was planned and executed under the guidance of Dr. Steven King, Director for Information Systems, Office of the Director of Defense Research and Engineering (DDR&E), representing the IATAC Steering Committee. The guiding philosophy, development methodology, and initial findings of the SOAR were presented to and vetted by members of the IATAC Steering Committee and other interested parties throughout the DoD IA and CND communities. The SOAR was developed in four phases, described below.

Phase 1: Analysis of source documents

The analysis team reviewed each source document and identified technical security-capability areas, solutions sets within them, and individual solutions within those solution sets, noting the status (current, planned, or needed) of each solution indicated by the source document.

Phase 2: Correlation of data into Capability Area Matrices, analysis to produce findings

For each capability area, the analysis team correlated, fused, and analyzed the single-document matrices' findings on solution sets and their component solutions. In the process of doing this, they organized the solutions and adapted solution definitions to fuse comparable solution descriptions in the different documents into a single, comprehensive description in the matrix. The resulting draft Capability Area Matrix plots each source document's position on every solution. (Was it mentioned at all? If so, what status was indicated?)

There are a few solutions that fall into more than one solution set and even fall into different capability areas. These include insider-accountability tools, which fall within both the scope of audit and insider-threat countermeasures, and appear in two different capability areas. Rather than attempting to cross-reference within or between matrices, the SOAR authors felt it would be more helpful to the reader to duplicate such solutions, including them within each Capability Area, solution set, and solution subset to which they pertained, and the repetition is noted within the solution description in the matrix.

After being assembled, the contents of the Capability Area Matrices were further analyzed, and a "roll-up" of status indicators were provided for one or more subsets of solutions within each solution set. For each document, the roll-up indicates the "farthest-out" solution in terms of status, documented by that source in that solution set. The purpose of the roll-up is (1) to provide the reader with a means of quickly determining whether a given source document includes any outstanding needs in a given solution set, (2) to compare across documents the presence of interest in a solution set, and (3) to determine any trends toward acknowledging outstanding needs across multiple documents.

Phase 3: Authorship of the report

Further analysis was performed to highlight significant trends (solution sets of interest vs. those not mentioned), areas of agreement or disagreement, and important omissions. An attempt was made to account for the stated purpose and intended audience of each source document to determine whether a given finding was, in fact, to be expected (i.e., consistent with the document's purpose and audience) or whether it was noteworthy given its stated purpose and audience. In cases where findings were consistent with a stated purpose and audience, further analysis was performed to determine whether an amendment to the document's purpose might be recommended, given the importance of the solution set and solution that is the subject of the finding.

Phase 4: Vetting of the SOAR

The SOAR outline, methodology, and initial findings were widely "socialized" across the organizations that make up the IATAC Steering Committee, including DDR&E, the Joint Staff J6I, the Defense Information Assurance Program (DIAP), the DISA Chief IA Executive, the National Security Agency (NSA) representing the US Strategic Command (STRATCOM) CND mission, and the DoD Enterprise-Wide solutions Steering Group (ESSG), representing the STRATCOM CND mission.

The SOAR and its findings were also briefed to the IATAC Leadership at the 2005 DoD IA Workshop and will be presented to the ASD NII.

The SOAR draft was reviewed by members of the IATAC Steering Committee. The final SOAR was then revised to ensure that it addresses all concerns and issues raised throughout the process. Final editing and production was then performed to produce the SOAR in hardcopy and electronic form; the latter was published in the IATAC repository.

Overview of SOAR findings and recommendations

Presentation of findings

The SOAR findings in each capability area are followed by the Capability Area Matrix to which they pertain. The explanatory notes and findings associated with and preceding each Capability Area Matrix are intended to clarify the scope of the capability area and to highlight those trends, disagreements, agreements, and omissions that are considered significant by the authors and, in some cases, to provide suggestions on how to address the issues documented in the findings. These findings report on the characteristics of solutions, the status (current, planned, or needed) of solutions, the relationships between solutions, terminology used to describe solutions, inconsistencies between findings, omissions from various source documents and the significance of those omissions, and other important observations.

The Capability Area Matrices are intended to give the reader an easy-to-digest depiction, for each individual solution, of the amount of agreement, disagreement, or lack of mention that exists across all source documents; for example, to quickly identify those documents that mention outstanding needs in each solution set and gain, across the documents, a collective sense of outstanding needs in each solution set. Further, the Matrices indicate, per document and per solution set, the "farthest-out" status of any solution in that solution set. Three status indicators are used in the Matrices to indicate the status of the solutions: [C] Current, [P] Planned, and [N] Needed.

Current solutions are those that are already deployed operationally or in a late-stage pilots or operational prototypes.

Planned solutions are those for which deployment or acquisition is planned or scheduled and budgeted. In practical terms, planned capabilities and solutions are expected to be deployed within the next five years; e.g., by the end of GIG Increment 1, 2008, or within the time frame of a current Program Objective Memorandum (POM), 2010.

Needed solutions reflect a need in the early stages of basic or applied research or recognized as an outstanding or unfunded need. Also, any solution in a long-term acquisition or deployment plan that would result in deployment after five years (e.g., GIG Increment 2 in 2012 or Increment 3 in 2016) is considered "needed" because it is planned beyond the period of any current DoD budget cycle.

The Capability Areas addressed in the SOAR are as follows:

IA Situational Awareness (IA SA)

IA SA includes all activities, technologies, policies, procedures, etc. required to develop a "picture" of the security posture of a system or network at any time as well as over time. Overlapping significantly with Attack Sensing & Warning (AS&W), the main elements of IA SA are generally agreed to include Detection and Sensing, Indication and Warning (I&W), Monitoring, Characterization, and Visualization. The SOAR authors also considered Network Mapping, Audit, and Threat Prediction to be part of IA SA.

CND Response

CND Response consists of initiating changes to existing protection systems or deploying additional protection measures in response to incidents detected and collected by sensors during AS&W. During CND response, the CND analyst interprets and acts on the IA SA data thus collected. In the SOAR, activities related to planning for CND response are also included in this capability area.

Network and Computing Infrastructure Protection

Sometimes also referred to as "network security," network and computing infrastructure protection comprises the Defense-in-Depth (DiD) protections of voice or data networks and the computing systems hosted on those networks against unauthorized access to or modification of information in transit over the network and against denial of service to or denial of access by authorized users.

Network and infrastructure protection includes those measures necessary to detect, document, and counter such threats. There are three categories of DiD for achieving network and infrastructure protection: (1) defense of the networks and their devices, (2) defense of the network boundary and external connections, and (3) defense of the networked computing infrastructure and environment.

Data Protection

Also referred to as "data security," the objective of Data Protection is to maintain the required levels of confidentiality, integrity, and availability of data, software, information services, applications, collaborative environments, etc. as they are accessed while "at rest" or "in transit." The most common threats to data include accidental or intentional modification, destruction, or disclosure to unauthorized users. Data protection in this SOAR includes assured information sharing.

Software Assurance and Application Security

Software Assurance and Application Security provides for continued safe, secure behavior of software applications in the face of an intentionally induced fault (such as that indicating an attack) and establishes software development life-cycle processes, test and assessment procedures, and technologies that ensure that the software applications produced are robust against attempted exploits and compromises. This is true whether the application is custom developed or assembled and integrated from software components acquired from third-party suppliers.

Availability

Availability is the assurance of timely, reliable access to data and information services for authorized users and systems. In a networked environment, this includes not only a user's ability to access hardware and software resources (such as user agents and servers) but also a user's ability to obtain a desired Quality of Service (QoS); e.g., to make use of network bandwidth with reasonable throughout. Network traffic must be able to traverse Local Area Networks (LANs) and Wide Area Networks (WANs), as required, to reach its intended destination.

Core Security Services

Core Security Services are a set of foundational security services comparable, but not identical, to the Primary Security Services described in the NSA's IA Technical Framework (IATF) and the Net-Centric Enterprise Services (NCES) Security Core Enterprise Services (CES). Core security services provide "security-enabling" technologies that support solutions across multiple capability areas, rather than being limited to a single capability area.

"IA Empowered" Workforce

The methods and technologies by which personnel gain the necessary level of understanding and awareness of the concepts and importance of IA necessary for them to use information systems, networks, and data in a way that does not expose these assets to avoidable compromises of their confidentiality, integrity, or availability.

The figure on the following page illustrates a portion of a Capability Area Matrix:

Following the Capability Area Matrices and Findings are a set of General Findings that pertain to issues or trends that cut across Capability Areas. Finally, the SOAR provides a series of observations and suggestions regarding DoD IA and CND strategy and planning, as characterized by the source documents.

The most significant findings that emerged from the analyses of the capability area matrices pertained less to gaps or overlaps than to omissions. Some important IA and CND solution sets and even capability areas that have received wide coverage in the press, have generated significant Information Assurance & Vulnerability Assessment (IAVA) activity, or have received much interest and even led to establishing programs in DoD were curiously overlooked by the DoD source documents. Similarly, some important security considerations related to technological initiatives in DoD were also overlooked by the source documents. The

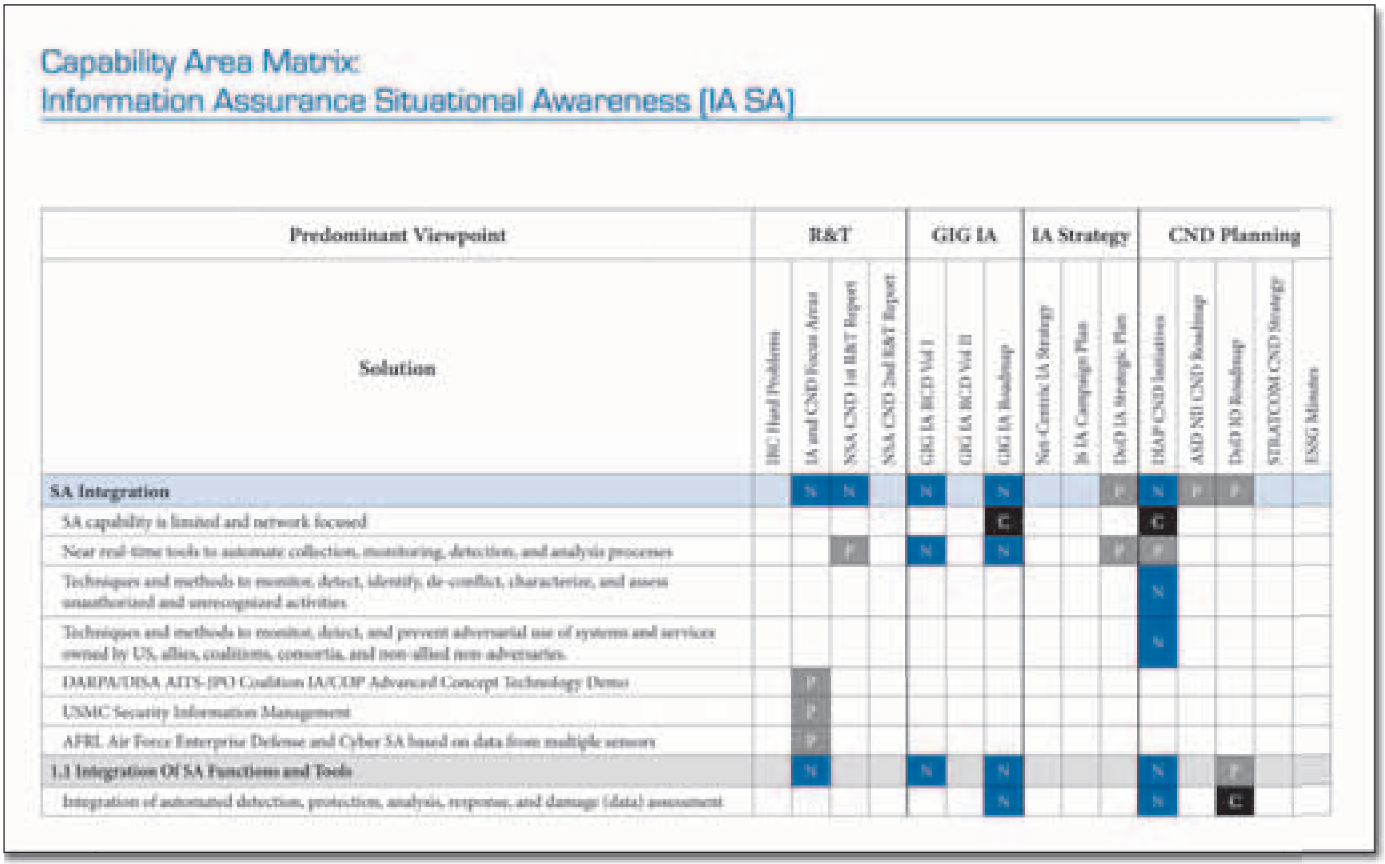

Capability Area Matrix:
Information Assurance Situational Awareness (IA SA)

Predominant Viewpoint	R&T				GIG IA			IA Strategy			CND Planning				
Solution	IRC Hard Problems	IA and CND Focus Areas	NSA CND 1st R&T Report	NSA CND 2nd R&T Report	GIG IA ICD Vol I	GIG IA ICD Vol II	GIG IA Roadmap	Net-Centric IA Strategy	J6 IA Campaign Plan	DoD IA Strategic Plan	DIAP CND Initiatives	ASD NII CND Roadmap	DoD ID Roadmap	STRATCOM CND Strategy	ESSG Minutes
SA Integration		N	N		N		N			P	N	P	P		
SA capability is limited and network focused							C				C				
Near real-time tools to automate collection, monitoring, detection, and analysis processes			P		N		N			P	P				
Techniques and methods to monitor, detect, identify, de-conflict, characterize, and assess unauthorized and unrecognized activities											N				
Techniques and methods to monitor, detect, and prevent adversarial use of systems and services owned by US, allies, coalitions, consortia, and non-allied non-adversaries											N				
DARPA/DISA AITS-JPO Coalition IA/CDP Advanced Concept Technology Demo		P													
USMC Security Information Management		P													
AFRL Air Force Enterprise Defense and Cyber SA based on data from multiple sensors		P													
1.1 Integration Of SA Functions and Tools		N			N		N				N		P		
Integration of automated detection, protection, analysis, response, and damage (data) assessment							N				N		C		

Figure 1: Sample Capability Area Matrix

SOAR reports on these omissions and on noteworthy gaps and overlaps among the source documents.

It was not possible for the analysts to determine, based on the stated purposes, objectives, and scope of each source document, the reasons why certain capabilities are discussed in some documents but not in other documents with similar stated purposes, objectives, and scopes. The SOAR analysts did not attempt to determine whether an omission (1) was predicated on the source document author's belief that the excluded capability had been or will be addressed by a different document, (2) reflected the source document author's belief that the capability was not desirable or important enough to be included, or (3) was simply overlooked by the author. Instead, solutions that appear to be important but were omitted from one or more given source documents were noted in the findings sections of the SOAR without speculation as to why each omission may have occurred.

A sampling of SOAR findings follows.

CND community and IA situational awareness

Better coordination may be required to ensure that the improved situational-awareness data content called for by the R&T and IA communities reaches the CND consumers more effectively, consistent with the CND planning call for better IA SA information sharing.

Relationship of audit to insider-misuse detection

It is critical that new insider-misuse detection tool dependencies on and interfaces to both existing and emerging audit capabilities be understood and acknowledged. This will ensure that capabilities provided by these applications are not unnecessarily duplicated across audit and insider-misuse detection tools, and that audit-data formats, collection, and reporting capabilities also satisfy the needs of the insider-misuse detection tools that rely on them.

Relationship of threat prediction to CND response planning

Threat predictions are a key input into CND response-planning activities. This said, the coverage of threat prediction is extremely scant across all source documents.

CND community apparently not driving automation of CND response

CND response-automation solutions appear predominantly in the R&T and GIG IA sources. Given their stated purposes, omitting needs related to CND response automation from the STRATCOM CND Strategy and ASD NII CND Roadmap is particularly noteworthy.

Need for attack attribution at the application level

To be comprehensive and effective, attribution techniques must address not only attacks against host and networks but also attacks on applications, which remain the most vulnerable components of Information Systems.

Planning for the impact of new technology adoption

The IA/CND Focus Areas document is the only source that identifies the need for planning to minimize the anticipated disruption caused by the transition to Internet Protocol Version 6 (IPv6) and other new, secure protocols.

Agreement on Ports, Protocols, and Services (PPS)

Not a single PPS solution is cited as an outstanding need by any of the four viewpoints. All source documents indicate that the PPS system is on its way to full deployment and usage.

Risks associated with Voice over Internet Protocol (VoIP) standardization

No viewpoint acknowledges the risks associated with DoD's planned migration to a Voice-over Internet Protocol (VoIP) network to the exclusion of other voice technologies. While redundancy is planned to provide backup VoIP connectivity in case of a successful attack on the network, the reduction in technical diversity increases the likelihood that a single-attack strategy could bring down both DoD's voice and data networks (both running on IP). If the attack vector is a worm or similar replicating agent, bringing backup VoIP networks online will only serve to provide additional targets to which the worm or self-replicating agent can spread.

Cross-domain solutions (CDS) for eXtensible Markup Language (XML)

If XML encryption is used in the GIG for XML messages or content that may at some point be required to cross a domain boundary, the CDS at that boundary will be unable to perform security decisions based on message content without "breaking" the Public Key Infrastructure (PKI) security model. (XML Encryption uses PKI.) A technique by which a Cross-Domain Solution (CDS) can perform content validation on public-key encrypted data (XML and other data formats) without violating the PKI's security model may be worth researching.

Federation of trust and cross-certification across domain boundaries

The ability to achieve assured information sharing and collaboration across domain boundaries will depend to a great extent on the ability to implement cross-domain certification and federation. This capability requires greater coverage across the source documents.

Failure to recognize the need for proactive software assurance and application security

Ninety percent or more of DoD's IA and CND capabilities are implemented in software, as are the same proportion of non-IA and CND capabilities those capabilities are intended to protect. With so much critical software forming the basis of DoD's net-centric enterprise and the GIG, Software Assurance and Application Security must be recognized as a critical Capability Area. Given the overwhelming concern across all viewpoints about the "insider threat", it is clearly time for both the IA and CND communities to recognize that the rogue developer and untrustworthy software supply chain represent two of the most insidious insider threats of all.

Privilege management for non-human actors

Current and planned solutions for Privilege Management focus exclusively on privileges of human users. Net-centricity is based on Web services and proxy agents that interoperate without human intervention. In the GIG, security for net-centric Web services will not be possible until the necessary privilege-management mechanisms are in place to enable establishing trust among Web services and proxy agents.

Apparently different definitions of "Community of Interest"

The term "Community of Interest" (COI) appears to be understood differently across different source documents. In GIG terms, a COI is defined primarily by its functional role (e.g., Focused Logistics vs. Training). In non-GIG sources, COI refers more often to a security-policy domain or enclave. This confusion in terminology makes it difficult to determine whether mandatory access controls or CDS will be required to enable information sharing and collaboration among COIs.

Coordination between CND reactivity and IA proactivity

The source documents do not indicate how much explicit coordination occurs between CND activities and IA activities to ensure that (1) they dovetail as seamlessly as possible by design, not by coincidence, to support each other, and (2) the information flows that must occur between them for each to be effective are not "broken" or throttled.

Of particular concern is whether a strong feedback loop exists between the CND impact and effectiveness of protection-assessment functions and IA solution specifiers and planners. This is necessary to ensure that the controls and protections that the CND community relies on as countermeasures and safeguards are as effective as they must be to minimize the impact of security incidents and to support rapid and full response and recovery. A similar feedback loop between Information Assurance Vulnerability Management (IAVM) and IA SA is also critical to ensuring that the IA SA "picture" is as complete and accurate as possible.

Challenges of implementing PKI in tactical and dynamically changing environments

There may be operational environments and situations, such as forward-deployed operations and rapidly changing and/or multi-national coalitions, in which use of a Public Key Infrastructure (PKI) may be challenging because of the following:

- difficulties associated with certificate and key distribution and management over tactical networks and infrastructure or in environments and operations with rapidly changing memberships, unpredictable infrastructures, and frequently changing security policies
- the inability to establish a single, shared trust root (i.e., certificate authority) or to federate trust roots in the multi-national and coalition environment and the inability to assure consistent connectivity to a root certificate authority in the tactical environment

It is important that the PKI Program recognizes these challenges and consider research into the suitability of a technology's Digital Rights Management (DRM), peer-to-peer encryption, and secure multicasting for use in these environments and situations.

Conclusion: Current SOAR and Next Steps

By highlighting the information assurance and cyber-security issues that are critical to our nation's defense, the findings of the IATAC SOAR should enable decision-makers in DoD to align the wide range of current IA and CND plans, roadmaps, and strategies, and to fill in the gaps that current strategies leaves untouched.

The findings should also prove helpful to the authors of future versions of the range of DoD IA and CND plans by highlighting the internal inconsistencies in current versions, as well as the inconsistencies between the various source documents.

The SOAR can also provide a basis for the different DoD organizations involved with IA and CND to align their knowledge and understanding of IA and CND terminology, concepts, and challenges. Specifically, the SOAR could provide a basis for developing a standard taxonomy of IA/CND capabilities to be used by authors of future versions of source documents, while the SOAR glossary could form a basis for a new standard DoD-wide glossary of key IA/CND terms.

Most importantly, the SOAR findings should also help DoD identify areas in which communication can be improved across the department's different IA/CND organizations - communication that should make DoD's IA and CND activities and capabilities more effective over the long term.

The IATAC team acknowledges that the inherent temporality of the source documents analyzed for the SOAR necessarily result in the inherent temporality of the SOAR's findings. The SOAR can only reflect a "snapshot" of DoD's IA and CND plans at the time those source documents were published.

To remain relevant, the SOAR needs to be updated periodically, to reflect new versions of current sources and determination if those sources are still relevant or if they have been superseded by new sources, as well as to incorporate findings from new sources.

Several of the SOAR's reviewers have also suggested that in future updates, the scope of the SOAR should be expanded, and a second level of analysis performed, in order to:

- Prioritize the findings according to their criticality;
- Verify/determine the reasons for observed gaps, overlaps, and omissions, and make specific recommendations for remediating them
- Address non-technology findings related to Tactics, Techniques and Procedures (TT&P; including Operations, Personnel, and Maintenance); Policy, Doctrine and Strategy; and Science and Technology (S&T);
- Acknowledging that DoD systems are also Federal systems, address overlaps and gaps among DoD and other Federal agency IA and cyber-security plans, particularly in the context of IA/CND capabilities that directly support the cooperation and collaboration between DoD's IA/CND organizations and those of other Federal agencies.

About the Author

Karen Mercedes Goertzel, CISSP

Karen Mercedes Goertzel, CISSP, has 24 years in IA, software assurance, and net-centric architectures and applications. She supports the Department of Homeland Security and ASD(NII) as a software assurance subject-matter expert, and was lead technologist for 3 years on DISA's Application Security Program. She also leads ASD(NII)'s effort to define an approach for two-level GIG network management.

Ms. Goertzel was lead author of DISA's 2002 IA Technology Forecast, Future IA Architecture, and IA Technology Roadmap. She has consulted in IA, software assurance, and CDS to DISA's Net-Centric Enterprise Services program, the National Geospatial Agency, NSA, the US/Canadian Anti-Drug Network, and Joint Staff, among others. In 2004, she authored IATAC Newsletter articles on Autonomic Computing and Computer Immunology.

Before joining Booz Allen, Ms. Goertzel was a CDS engineer for Getronics (now BAE/DigitalNet). For Wang and Honeywell she consulted to DoD, Civil agencies, and NATO on security architecture and policy, risk management, COOP, and software development.

IATAC Spotlight on Research

Dartmouth College

by Ronald Ritchey

This is the third article in a series that spotlights important activities in Information Assurance (IA) education and research and describes the latest projects at some of the nation's best IA academic centers.

The featured center for this issue is the Institute for Security Technology Studies (ISTS), located at Dartmouth College, Hanover, NH. [1] ISTS was founded by the US Congress in 2000 to perform counterterrorism research. Its mission is to strengthen homeland security through research, education, and outreach programs that focus on technology that is critical to cyber security, emergency preparedness, and response. ISTS has two main centers: the Cyber Security and Trust Research Center (CSTR) and the Emergency Readiness and Response Research Center (ER3C).

CSTR [2] concentrates on cyber-security issues. Its focus covers both reactive and preventative technologies and includes research into methods that can be used to protect computer systems against attack, secure communications between computer systems, and improve detection of malicious computing activities. The Center also addresses the affect of social, political, and economic considerations on the security of computer systems.

One current project at CSTR is the study of Process Query Systems (PQS). These systems examine how to extract processes from an observed event stream, based on how closely the observed events match a description of the process' state. The underlying model assumes that, in any system, there are observable events caused by processes within the system. The events that occur are directly related to the internal state of these processes, and, by observing the events, at least some of these internal states can be inferred. The PQS can observe a sequence of events (e.g., event stream) and determine which processes produced which events. This is very useful in a broad variety of detection and prediction applications, including intrusion-detection systems. ISTS has applied its PQS research across such diverse domains as network- attack prediction (including the early detection of worms), Supervisory Control and Data Acquisition (SCADA) systems, and fish-tracking. [3]

Another area of active research at CSTR is Public Key Infrastructure (PKI). ISTS hosts the PKI/Trust lab, a leading PKI research facility in academia. [4] One key problem that prevents wide PKI deployment is the difficulty in building the infrastructure necessary to securely associate keys with people. ISTS researchers are working to show how to build good PKI infrastructures and are learning how to do this, in part, by applying their findings broadly within Dartmouth. For instance, students can gain access to campus computer systems using the lab's PKI. The lab also hosts the Higher-Ed Bridge Certificate Authority (HEBCA), which will be used across many academic institutions to federate PKI certificates. When fully deployed, it will be the world's largest bridge Certificate Authority.

The second research center, the ER3C, [5] examines how to use information technology to improve our response to large-scale disasters. This research is focused on the needs of first responders and draws from a number of applied research sources to develop new solutions to problems in emergency preparedness and response. One area of major focus is situational awareness—a key requirement during a disaster. Another area is training and simulation. ER3C has provided several interesting simulations and exercises that have provided emergency responders with insight into the best methods to reduce the impact of specific emergency situations. In one recent example, teams and technology from the ER3C supported emergency responders in Lebanon, NH, as they completed an annual tabletop exercise of disaster response at the local airport.

ISTS is also a member of the Institute for Information Infrastructure Protection (I3P), which is managed by Dartmouth. [6] The I3P is a consortium comprising 24 organizations drawn from academia, federally funded labs, and non-profit organizations. Its goal is to bring experts together to identify and help mitigate threats aimed at the US information infrastructure. Operational since 2002, it

continued on page 16…

The Kerf Toolkit for Intrusion Analysis

by Javed Aslam, Sergey Bratus, David Kotz, Ron Peterson, and Daniela Rus

Network-based intrusions have become a significant security concern for system administrators everywhere. Existing Intrusion Detection Systems (IDSs), whether based on signatures or statistical learning of normal behavior, give too many false positives, miss intrusion incidents, and are difficult to keep current with all known attacks. Although recent high-level correlation tools have improved the quality of alerts to system administrators [1], [2], IDSs have a limited success rate, tend to detect only known attack types, and ultimately result in only an alert message to a human administrator. (*In this paper, we will not discuss the relatively recent development of Intrusion Prevention Systems, that offer active response to an intrusion without human intervention*). Thus human experts are still required to analyze the alert (and related data) to determine the attack's exact nature. Human experts are also the key tool for identifying, tracking, and disabling new attack forms. This work often involves experts from several organizations working together to share their observations, hypotheses, and attack signatures. Unfortunately, few tools help these experts in the process of analyzing log data.

To alleviate this situation, we developed the Kerf toolkit (so named for a kerf, which is the slit made by a saw as it cuts through a log). Its goal is to provide an integrated set of tools that aid system administrators in analyzing the nature and extent of an attack and then communicating the results to other administrators or law-enforcement agencies. Kerf contains semi-automated tools that help system administrators identify attack characteristics based on data from network and host-based sensors, develop a hypothesis about an attack's nature and origin, express and share that hypothesis with security managers from other sites (without sharing actual log data, which may be sensitive for their organization), test the hypothesis at other sites, and coordinate the testing results.

Kerf and intrusion analysis

Picture the typical System Administrator, responsible for a collection of hosts on one or several organizational subnets. Each host logs its activity using the Unix syslog facility or the Windows Event Logging service. An IDS monitors some or all hosts—possibly the entire network—and generates and logs alerts about potential attacks. Once a system administrator discovers an attack, he or she must put on an analyst hat and further investigate (see Figure 1).

Kerf is intended to assist in this investigation, commonly referred to as intrusion analysis, after an attack is detected. We assume that correct and complete host and network logs are available, up to a point. To ensure this, Kerf includes agents installed in monitored machines that forward encrypted log records to a secure, off-host logging server (see Figure 2). The analyst goal, then, is to reconstruct evidence of an attack from individual event records in the available logs.

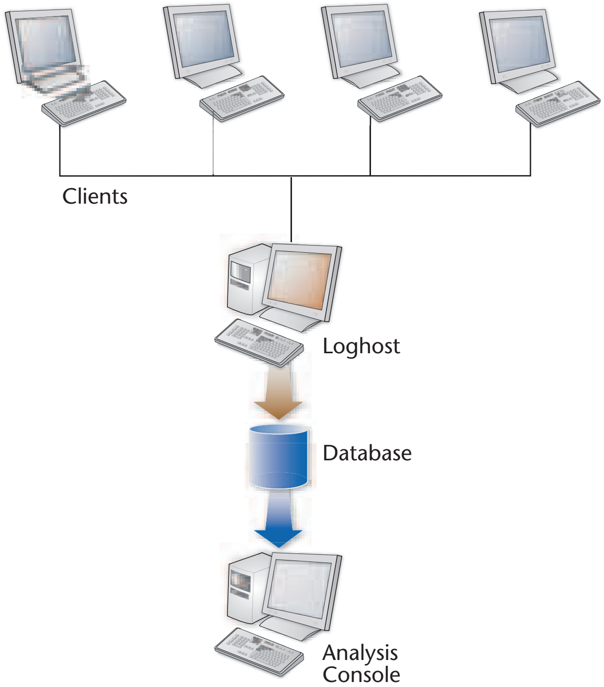

Figure 1: Overview of Kerf physical architecture

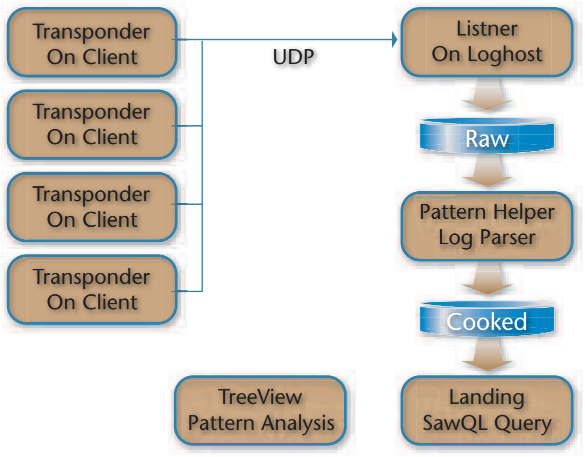

Figure 2: Overview of Kerf software architecture

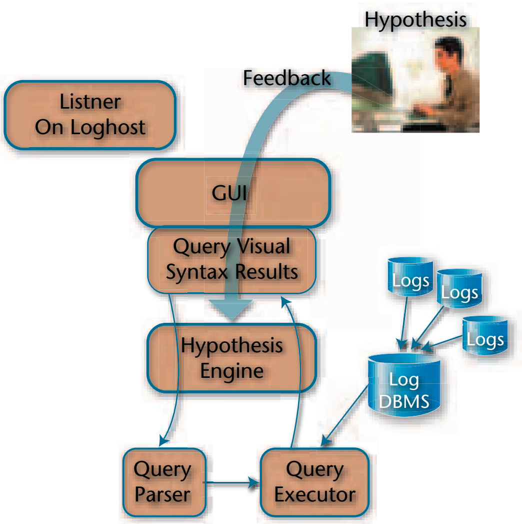

Figure 3: Hypothesis refinement: the Kerf approach

The analysis process is inherently interactive: an analyst begins with a vague mental hypothesis about what happened and then uses Kerf tools to test and revise that hypothesis (see Figure 3).

The process is also inherently iterative: each new piece of information permits the analyst to revise the hypothesis and explore further. The hypothesis is refined, as information that partially confirms it is discovered, and is expanded, as the analyst tries new approaches that broaden the investigation. The result is a specific hypothesis about an attack's source and nature and the concrete evidence to support the hypothesis.

Many tools for parsing text-based system logs currently available to system administrators [3], [4], [5] rely on extensions of *regular expressions*, which require syntactically complex constructions to search logs for relevant entries or to extract relevant parameters from them. This, in turn, often requires writing ad hoc scripts to correlate events from different logs or hosts. A number of tools that store parsed logs in relational databases, such as the Microsoft LogParser [6] (for which Burnett [7] is an excellent tutorial), permit users to express certain correlations in the form of Structured Query Language (SQL) queries with joins, but such expressions very quickly tend to grow intractable. In such conditions, any systematic recording of hypotheses, actions, and results for later study becomes very difficult. Because the analysis process is difficult and tedious, most system administrators can't fully explore and understand an attack or document it so that others can study it. Kerf aims to make intrusion analysis more efficient by providing the following:

- A secure mechanism for network and host logging to a dedicated log server, which keeps the logs' records in a relational database

- A correlation engine that accepts queries in SawQL, a domain-specific extension of SQL that is designed to concisely describe sequences of records correlated on their various parameters including their timestamps

- The PatternHelper tool to help a user write patterns for extracting parameters from free text-log formats, such as UNIX syslog
- A User Interface (UI) front end, Landing, that adaptively organizes large result sets from SawQL queries for more convenient viewing and analysis and permits a user to attach his or her own tags to records; in particular, to mark them as "suspicious" or "innocuous"
- The Hypothesis engine (under development) that aids a user with query generalization and refinement by learning from user feedback and adjusting the application's data organization algorithms or by suggesting new queries

Event correlation

It is natural to describe an intrusion as a sequence of events, some of which leave their traces in the form of records in various logs. These records are likely to be correlated on some parameters (e.g., the corresponding events may originate from or take place on the same host or may involve a logged common value associated with some protocol). Even more likely, they will be correlated on time (e.g., one event occurs before or after another, within a short period of time).

SawQL, the SQL-based Kerf query language, permits convenient expression of relative or absolute temporal and parameter correlations at the same time that it abstracts away the gory details of database joins. Thus queries in SawQL naturally represent sequences of correlated events and can be used to express and share hypotheses. Examples of SawQL expressions that describe actual intrusions can be found on the Kerf project Web site. [8]

Data organization and presentation

In the practice of intrusion analysis, there inevitably occurs a scenario in which a query returns many screenfuls of matching log records; each of which are full of diverse records; refining the query appears possible only after the majority of these records have been examined. In such situations, automated data organization algorithms that attempt to summarize and classify the data can save an analyst time and effort. Kerf uses entropy-based, recursive data organization algorithms to produce a tree-form representation of query results every time the results exceed a user-defined threshold size.

More precisely, the records are grouped by the unique values of their parameters. The order of grouping is chosen adaptively, based on the distributions of values of each parameter across the given result set. The resulting groups correspond to intermediate nodes of the tree, which are marked with the parameter values common to all records contained under a node. Thus the upper levels of the tree serve as a summarization of the result set.

An important side effect of this grouping method is that it will likely highlight "abnormal" events, which are of greatest interest in attack analysis. The data organization algorithm is tuned to produce trees of moderate depths and branching factors to aid the following typical tasks:

- Discovering the actual composition of result sets
- Understanding the distribution and ranges of selected parameter values and finding subsets of records with anomalous values
- Navigating to subsets of interest
- Extracting subsets of interest for use with another query

The snapshot in Figure 4 below shows a set of 1357 Snort portscan alerts, grouped first by destination port and then by source and destination IPs in Frame (A). Frame (B) summarizes the value ranges of other parameters in the selected group. Both Frames (A) and (B) can be used for user tagging of groups or individual records (not shown). Frame (C) accepts commands in an internal scripting language, and Frames (D) and (E) show status messages.

A user can add levels of grouping or define his or her own classification tree templates, bypassing the algorithm entirely or running it only on subtrees of a pre-defined classification. This method is useful in cases when the overall expected structure of the log data is well understood, whereas seeing where a new batch of records ends up in a pre-defined classification may provide a useful clue. All user operations on a dataset can be recorded and replayed on other comparable result sets. A user can also directly define his grouping and classification rules in an internal template language.

Kerf users will notice that the simplest operations on group nodes of a tree (i.e., subsets of the result set) are functionally similar to UNIX command chains—*"grep ... | sort | uniq -c | sort -n"* or *"select distinct ... group by ... order by ..."* statements of SQL environments—while providing much more flexibility in defining and connecting the filters and in keeping all records within a common and reusable classification framework.

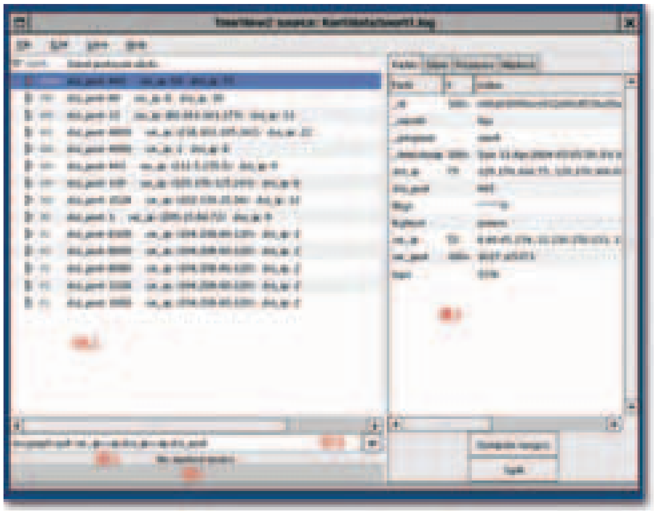

Figure 4: Application snapshot

An example of adaptive data organization

The following example shows how adaptive data organization can elucidate the structure of a moderately sized result set at a single glance. Here, a flat list of authentication records from an actual UNIX system log, selected by a simple query without correlation, is presented as an adaptively constructed tree. The user is a System Administrator who is concerned with logins from the network of a certain Internet Service Provider (ISP) and wants a brief summary of failed and successful logins. A query for login events from *.isp.net returns some 600 records.

Subsequently, it is determined that all logins originated from two legitimate users who happened to inhabit distinct dynamic IP ranges, one of whom was prone to typos. The feature pair (user, host) was found by the data-organization algorithm to produce the best tree form. The user was thus presented with a 12-line summarization of the 600-line result set. It also became clear that most logins came from one user and his login records were further grouped by month. (See Figure 5.)

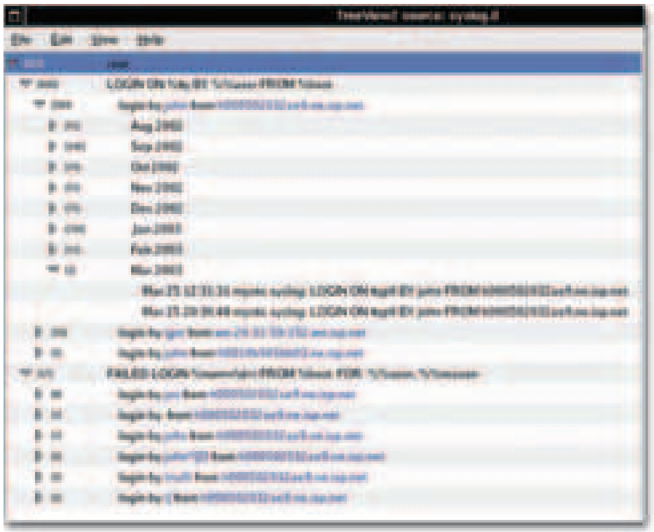

Figure 5. Treeview, some nodes expanded. The Kerf module for adaptive display of query results chose this summarization of the 600+ login events from *.isp.net.

Work in progress

In the near term, we plan to extend our system to handle other types of logs; in particular, IDS logs in the Intrusion Detection Message Exchange Format (IDMEF) and kernel audit logs, such as Sun Solaris BSM [9] and Linux Snare [10] and Syscalltrack. [11]

In the long term, a major goal of the Kerf project is to provide semi-automated tools to aid an analyst in hypothesis generation, refinement, archiving, generalization, and extrapolation. To this end, we are developing the following:

- A hypothesis engine, consisting of a hypothesis-generation module to assist a user in formulating the initial hypothesis
- A hypothesis-refinement module to assist in modifying the initial hypothesis to better target suspicious behavior
- A hypothesis-sharing module to assist in taking the final hypothesis and archiving it for later use, extrapolating it for other specific users and domains, and generalizing it for wider applicability

We expect our new algorithms and tools to be a unique contribution to the current state of intrusion analysis, by automating the existing best-of-breed analysis practices, and offering new powerful and flexible data organization techniques.

Acknowledgments

Marco Cremonini and Andrea Schiavoni designed and implemented an earlier prototype of a similar log-collection and analysis tool. We thank Giovanni Vigna and George Bakos for their helpful discussions and their assistance in locating host and network-log data. We also appreciate the input of other members of Dartmouth College's Institute for Security Technology Studies (ISTS), who commented on the current state of the art and the needs of IDS tools. Kerf is a project of the Institute for Security Technology Studies at Dartmouth College, supported under Award number 2000-DT-CX-K001 from the U.S. Department of Homeland Security, Science and Technology Directorate. Points of view in this document are those of the author(s) and do not necessarily represent the official position of the U.S. Department of Homeland Security or the Science and Technology Directorate. ■

References

[1] Haines, J., Dorene Kewley Ryder, Laura Tinnel & Stephen Taylor. (2003, January/February). Validation of sensor alert correlators. IEEE Security & Privacy, 1(1), 46–56.

[2] Peng N, Yun Cui & Douglas S. Reeves. (2002). Analyzing intensive intrusion alerts via correlation. In Proceedings of the Fifth International Symposium on Recent Advances in Intrusion Detection (RAID), Vol. 2516 of Lecture Notes in Computer Science, pp. 74–94. SpringerVerlag, October 2002. [Online] Available: http://link.springerny.com/link/service/series/0558/papers/2516/25160074.pdf

[3] Bird, T. & Ranum, M. Log analysis resources. http://www.loganalysis.org/

[4] Chuvakin, A. (August 2002). Advanced log processing. SecurityFocus.com [Online]. Available: http://online.securityfocus.com/infocus/1613

[5] Allison, Jared. (2002) Automated log processing.; login:, 27(6), 17–20.

[6] http://www.microsoft.com/windows2000/downloads/tools/logparser/

[7] Burnett, Mark (2003). Forensic log parsing with Microsoft's logparser. Loganlaysis.0rg [Online]. Available: http://www.securityfocus.com/infocus/1712, July 2003 7.

[8] http://kerf.cs.dartmouth.edu

[9] http://www.sun.com/software/security/audit/

[10] www.intersect alliance.com/projects/Snare/

[11] http://syscalltrack.sourceforge.net/

About the Authors

Javed Aslam

Javed Aslam is an associate professor in the College of Computer and Information Science at Northeastern University. His research interests include machine learning, information retrieval, computer security, and algorithm design and analysis. Aslam received a BS in electrical engineering and mathematics from the University of Notre Dame and a PhD in computer science from the Massachusetts Institute of Technology (MIT). He is a member of the Association for Computing Machinery (ACM) and the Institute of Electrical & Electronics Engineers (IEEE). He may be reached at jaa@ccs.neu.edu.

Sergey Bratus

Sergey Bratus is a postdoctoral research associate in Dartmouth College's Computer Science Department. His research focuses on applying machine learning and Artificial Intelligence (AI) techniques to intrusion analysis. He received his undergraduate education at the Moscow Institute of Physics and Technology and his PhD from Northeastern University. He is a member of Usenix and an associate member of the Free Software Foundation. He may be reached at sergey@cs.dartmouth.edu.

David Kotz

David Kotz is a professor of computer science at Dartmouth College, director of the Center for Mobile Computing, and executive director of the ISTS. His research interests include context-aware mobile computing, pervasive computing, wireless networks, and intrusion detection. He received an AB in computer science and physics from Dartmouth and a PhD in computer science from Duke University. He is a member of the ACM, IEEE Computer Society, Usenix, and Computer Professionals for Social Responsibility (CPSR). He may be reached at dfk@cs.dartmouth.edu.

Ron Peterson

Ronald Peterson is a senior programmer in Dartmouth College's Computer Science Department and owner of Peterson Enterprises, which develops PC-based, Musical Instrument Digital Interface (MIDI) musical instruments and graphics software. His research interests include cybersecurity, wireless sensor systems, cattle herding, mobile agents, and machine-vision interfaces for novel musical instruments. He received a BA in physics from Lawrence University. He may be reached at rapjr@cs.dartmouth.edu.

Daniela Rus

Daniela Rus is an associate professor in the Electrical Engineering and Computer Science department at MIT. Her research interests include distributed robotics, mobile computing, and self-organization. She has a PhD in computer science from Cornell University. She was the recipient of an National Science Foundation (NSF) Career Award, is an Alfred P. Sloan Foundation Fellow, and a Class of 2002 MacArthur Fellow. She may be reached at rus@csail.mit.edu.

continued from page 11…

"IATAC Spotlight on Research—Dartmouth"

has been engaged in identifying and addressing critical research areas required in cyber security and critical-infrastructure protection. One result of its efforts is the I3P cyber security Research and Development (R&D) agenda, which identifies critical gaps in cyber security and provides a list of recommended research priorities. [7]

The I3P Consortium recently launched two major cyber-security research projects that involve half the I3P's member institutions. Over the next two years, research teams will focus on developing models, tools, and technologies to protect SCADA systems used in the oil and gas industry and to gain a better understanding of the economic factors influencing cyber-security decisions. The first project, launched in March 2005 and led by Sandia National Laboratories, is an $8.5M effort to identify SCADA vulnerabilities and the interdependencies between SCADA systems and other critical infrastructures. [8] Researchers will develop metrics and models for assessing and managing SCADA security and will create next-generation SCADA systems with built-in security. The second research initiative, led by the RAND Corporation and worth $3M over two years, will help quantify the costs of cyber attacks and measure the effectiveness of current security tools and policies. [9]

I3P also supports a fellowship program designed to increase the number of cyber-security experts and researchers to fill the gap areas it has identified. [10] This program provides up to $150,000 in financial support for successful applicants. Five fellows are appointed each year, and the fellows are required to conduct research at one of the I3P member organizations. To be eligible for the program, research candidates must have received their doctorate no more than three years ago and have strong backgrounds in fields related to the gap areas. While the 2005 fellows have already been determined, a call for proposals will be released later this year for the 2006 program. ■

References

[1] http://www.ists.dartmouth.edu/
[2] http://www.ists.dartmouth.edu/cstrc/mission.php
[3] http://www.pqsnet.net/projects.php
[4] http://www.dartmouth.edu/%7Epkilab/
[5] http://www.ists.dartmouth.edu/er3c/mission.php
[6] http://www.thei3p.org/
[7] http://www.thei3p.org/about/2003_Cyber_Security_RD_Agenda.pdf
[8] http://www.thei3p.org/about/news/20050302_scada.html
[9] http://www.thei3p.org/about/news/20050516_econ.html
[10] http://www.thei3p.org/fellowships/index.html

IATAC Spotlight on Subject Matter Expert (SME)— Dr. Sergey Bratus

by Ronald Ritchey

This article is the third in a series of profiles of members of the Information Assurance Technology Analysis Center (IATAC) Subject Matter Expert (SME) program.

This article is the third in a series of profiles of members of the Information Assurance Technology Analysis Center (IATAC) Subject Matter Expert (SME) program. Information Assurance (IA) and Information Operations (IO) experts from many different types of organizations volunteer to be IATAC SMEs and provide information on their areas of expertise, education and training, professional certifications, inventions, and patents. When Department of Defense (DoD) or other government personnel contact IATAC with questions regarding IA or IO, IATAC can leverage its SME database to identify members who are particularly well suited to answering those questions. SMEs are also encouraged to contribute papers and other materials to IATAC's Scientific and Technical Information (STI) collection. The work of the SMEs furthers our understanding and capabilities in IA.

The IATAC SME profiled in this article is Dr. Sergey Bratus, a Postdoctoral Research Associate at the Computer Science Department at Dartmouth College, Hanover, NH. [1] Dr. Bratus is one of the researchers working on the Kerf project (highlighted in another article in this issue). Kerf provides system administrators with a unique collection of tools to analyze logs and alerts to identify intrusions into their networks. [2] Dr. Bratus's background is in Mathematics, which has provided a foundation for much of his research. He is currently working on machine learning and other methods of artificial intelligence to automate intrusion analysis. Using approaches based on information theory that have proved efficient in other areas, such as natural language processing, the Kerf group is developing flexible approaches that are less dependent on data format than those used in current signature-based methods.

Kerf uses data-organization algorithms to present the logged data in such a way that its statistical anomalies become apparent to the operator. Initially, the algorithm takes an unsupervised guess to choose the best data presentations; then it uses operator feedback to modify its data views and to adjust its concepts of anomalous vs. normal input. This is an interactive, iterative process that can be used to adjust what is considered anomalous as, over time, attack or normal behavior patterns shift.

The prototypes of Kerf tools are already in use at Dartmouth, where they are helping local administrators to gain valuable insights into their network data. However, the project could benefit from additional data and use cases. If you are interested in getting involved in the project, IATAC can assist in putting you in touch with Dr. Bratus.

Dr. Bratus has a keen interest in the low-level details of cyber attacks and kernel-level Operating System (OS) countermeasures. His initial introduction to intrusion analysis occurred when the Linux servers he administrated came under attack. Seeing first-hand the impact of these attacks has fueled his interest in how attacks work and also in methods that can be used to recognize and analyze attacks.

Dr. Bratus has a broad interest in UNIX security, including security operations within the kernel. He is currently working on a project, directed by Dr. Doug McIlroy and Dr. Sean Smith, to introduce structure to SELinux security policies to make them more manageable for system administrators and more amenable to automated verification. He is also interested in reverse-engineering malware and enjoys collecting Linux kernel rootkits.

If you have a technical question for Dr. Bratus or other IATAC SMEs, please contact iatac@dtic.mil. The IATAC staff will assist you in reaching the SME best suited to helping you solve the challenge at hand. If you have any questions about the SME program, or if you are interested in joining the SME database and providing technical support to others in your domains of expertise, please contact iatac@dtic.mil, and the URL for the SME application will be sent to you. ■

References

[1] http://www.cs.dartmouth.edu/

[2] http://kerf.cs.dartmouth.edu/

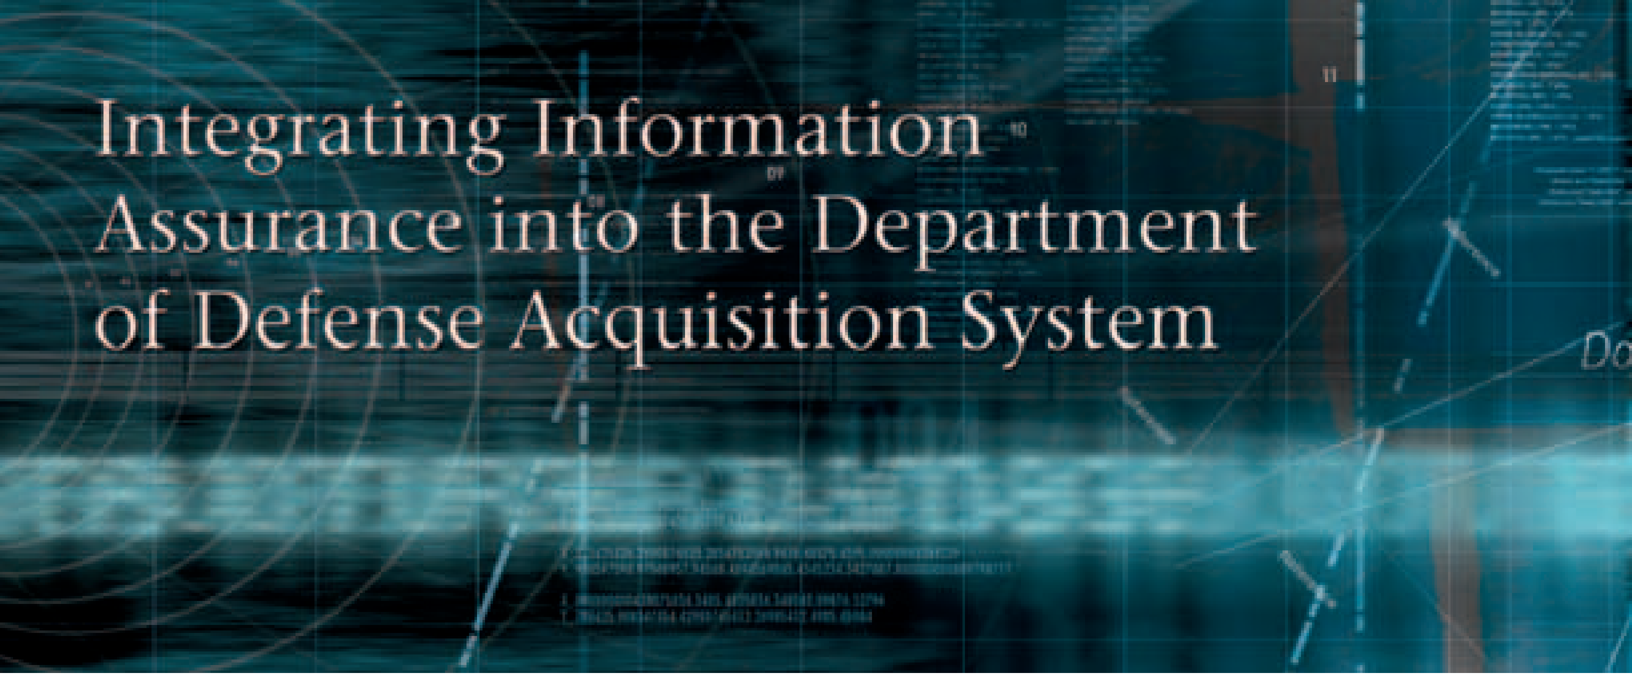

Integrating Information Assurance into the Department of Defense Acquisition System

by Eustace King, Arthur King, and Dominic Cussatt

Most Department of Defense (DoD) acquisition programs that deliver capability to warfighter or business domains will use Information Technology (IT) to enable or deliver that capability. For these programs, developing a comprehensive and effective approach to implementing Information Assurance (IA) is a fundamental requirement and will be a key to successfully achieving program objectives. IA is defined as "measures that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for the restoration of information systems by incorporating protection, detection, and reaction capabilities."

There is no place at which IA is more important than in the acquisition process. Planning for and implementing appropriate IA practices during the acquisition process will ensure that IA is "baked in" to the system rather than "brushed on" afterwards resulting in a secure and "Net-ready" system with less impact on cost and schedule. The information that follows identifies recently developed policies from which IA requirements are derived, addresses how IA is implemented in acquisition documentation, and lists several sources of IA assistance and supporting resources.

New policy and guidance

DoDI 8580.1, *Information Assurance in The Defense Acquisition System*

The intent of Department of Defense Instruction (DoDI) 8580.1, Information Assurance in the Defense Acquisition System, is to make existing IA policy more accessible and more easily understood by the acquisition community. It describes required and recommended levels of IA activities as they pertain to the acquisition of systems and services. It also describes the essential elements of an Acquisition IA Strategy and its applicability and prescribes an Acquisition IA Strategy submission and review process. DoDI 8580.1 captures the acquisition-related IA policies of DoDD 8500.1 and DoDI 8500.2. It also considers the companion implementation guidance developed jointly by the Assistant Secretary of Defense for Networks and Information Integration (ASD NII) and the Office of the Under Secretary of Defense (USD) for Acquisition Technology & Logistics (AT&L) and published as the Defense Acquisition Guidebook (DAG). The DAG provides specific "how-to" information required by acquisition programs. Major points within DoDI 8580.1 are as follows:

- All acquisitions of mission-critical or mission-essential IT systems, as defined in DoD Instruction 5000.2, Operation of the Defense Acquisition System, May 12, 2003, shall have an adequate and appropriate Acquisition IA Strategy that shall be reviewed prior to all acquisition-milestone decisions, program-decision reviews, and acquisition-contract awards.
- Heads of DoD Components will ensure that IA is implemented
- Program Managers (PMs) will perform the following:
  - Appoint an IA Manager
  - Ensure Mission Assurance Category (MAC) and confidentiality levels are identified
  - Identify baseline IA controls
  - Integrate IA
  - Plan and execute IA Certification & Accreditation (C&A)
  - Provide updates to Integrating Integrated Product Team (IIPT) and Overarching Integrated Product Team (OIPT)
- Acquisition IA Strategy will be approved by the appropriate Component's Chief Information Officer (CIO) and reviewed by the DoD CIO for all Major Automated Information System (MAIS) and Acquisition Category (ACAT) 1D programs
- DoD CIO review of Acquisition IA Strategies for all other programs is delegated to the Component CIO

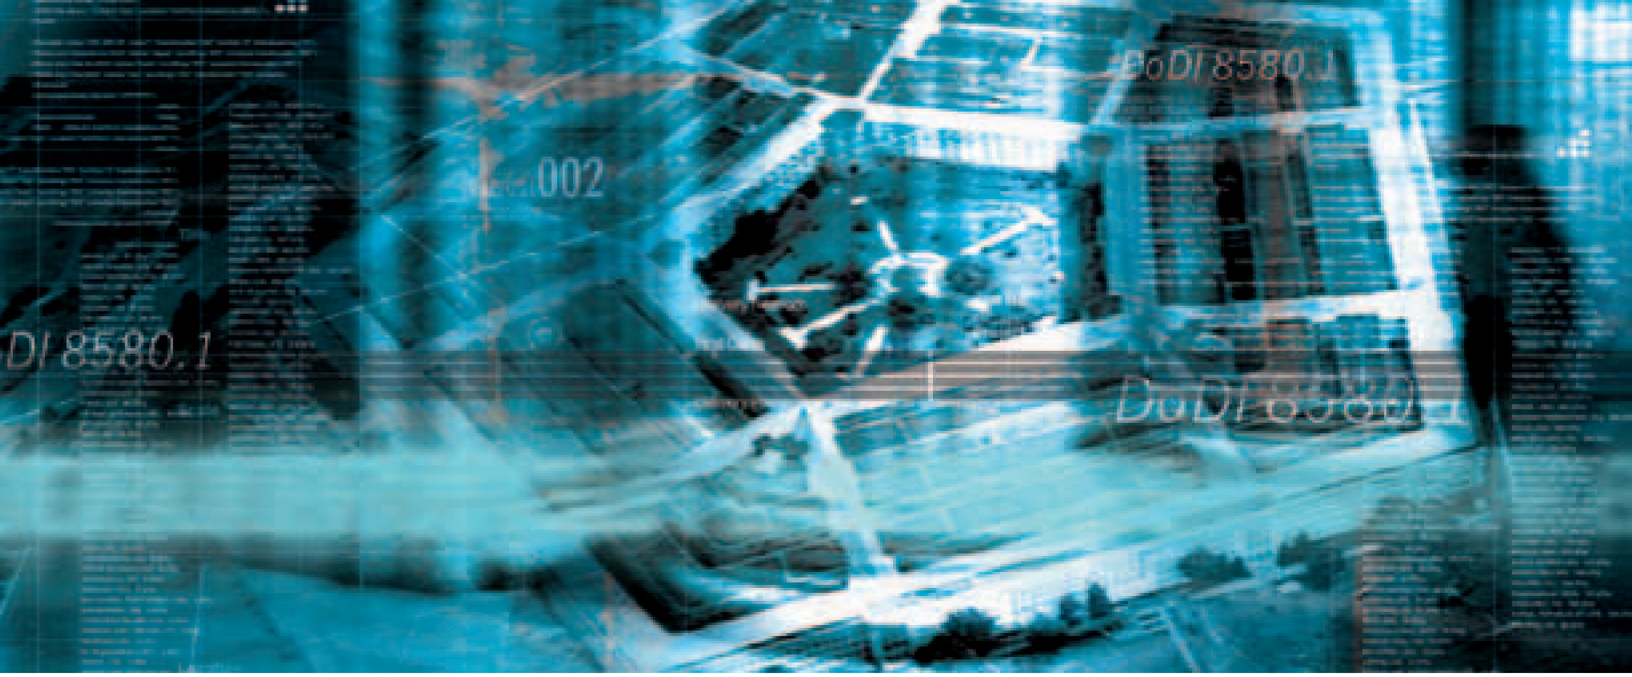

Figure 1 on the following page illustrates the relationships among DoDI 8580.1 and other regulations.

IA Guidance in the *Defense Acquisition Guidebook (DAG)*

PMs must be familiar with and understand the following guidelines and processes:

- Statutory and regulatory requirements governing IA and the major tasks involved in developing an IA organization
- Defining IA requirements
- Incorporating IA into a program's architecture
- Developing an acquisition IA strategy (when required)
- Conducting appropriate IA testing
- Achieving IA certification and accreditation for a program

DoD policy and implementing instructions on IA can be found in DoD's 8500 series publications.

The ASD NII recognized that the principle IA policy documents, DoDD 8500.1 and DoDI 8500.2, were written for IT and IA professionals and were not "user friendly" for acquisition professionals. The essential elements of these documents were extracted, sequenced to fit the acquisition life cycle, and expressed in terms that are relevant to system-acquisition professionals. The IA section of the DAG identifies requirements for IA compliance, the policy from which they are derived, their relationship to the acquisition framework, and the details that should be considered in developing an effective Defense-in-Depth (DiD) IA approach in a Net-centric environment.

The DAG is located online at http://akss.dau.mil/DAG/welsome.asp. The IA section can be found in Chapter 7.5 and contains the following:

- IA overview
- Mandatory policies
- IA integration into the acquisition life cycle
- Estimated IA activity durations and preparation lead times
- Integrating IA into the acquisition process
- PM responsibilities
- IA controls
- IA testing
- Acquisition IA strategy
- DoD Information Technology Security C&A process (DITSCAP)
- Software security considerations
- IA definitions
- IA considerations for the TEMP, SEP, and Acquistion Strategy

Other emerging guidance

CJCSI 3170.01, Joint Capabilities Integration and Development System (JCIDS)

This document establishes the policies and procedures of the JCIDS. These procedures support the Chairman of the Joint Chiefs of Staff (CJCS) and the Joint Requirements Oversight Council (JROC) in identifying, assessing, and prioritizing joint military capability needs. It also calls for transitioning to Net-Ready Key Performance Parameters (KPP), which include IA.

Addressing IA in acquisition documentation

Successfully implementing IA into a system that is being acquired requires a program's IA approach to be visible, coherent, and supportable across all program strategies, plans, and activities. Key acquisition documents

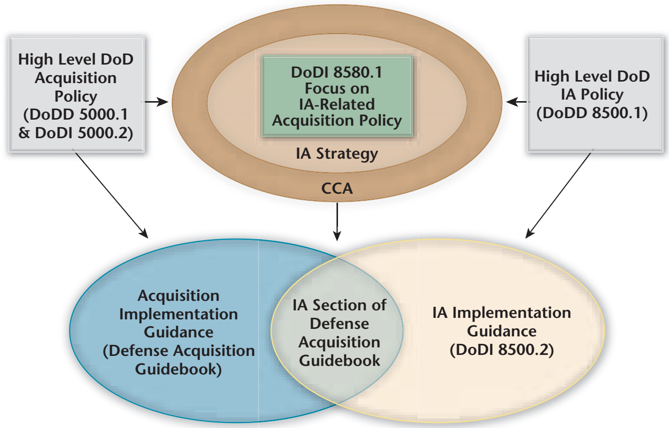
High Level DoD Acquisition Policy (DoDD 5000.1 & DoDI 5000.2)
DoDI 8580.1 Focus on IA-Related Acquisition Policy
IA Strategy
CCA
High Level DoD IA Policy (DoDD 8500.1)
Acquisition Implementation Guidance (Defense Acquisition Guidebook)
IA Section of Defense Acquisition Guidebook
IA Implementation Guidance (DoDI 8500.2)

Figure 1. DoDI 8580.1 as it relates to other statutory and regulatory guidance

that must be synchronized in IA include the Acquisition IA Strategy, the Acquisition Strategy, and the Test and Evaluation Master Plan (TEMP). These documents are discussed in the following sections, which provide an overview of each and identifies their key IA considerations. Detailed information can be found in the DAG.

Acquisition IA Strategy

The primary purpose of the Acquisition IA Strategy is to ensure compliance with the statutory requirements of the Clinger—Cohen Act and related legislation, as implemented by DoD Instruction 5000.2. As stated in that Instruction, the Acquisition IA Strategy provides documentation that "The program has an information assurance strategy that is consistent with DoD policies, standards, and architectures, to include relevant standards." A PM develops an Acquisition IA Strategy to help the program office organize and coordinate its approach to identifying and satisfying IA requirements consistent with DoD policies, standards, and architectures.

Developed early in the acquisition life cycle and written at a high level, the Acquisition IA Strategy documents a program's overall IA requirements and approach, including the C&A approach. The Acquisition IA Strategy lays the groundwork for a successful C&A process by facilitating consensus among the PM, Component CIO, and the DoD's CIO on such pivotal issues as:

- determining the Mission Assurance Category, the Confidentiality Level, and applicable Baseline IA Controls;
- selecting an appropriate C&A process;
- identifying a Designated Approving Authority (DAA) and Certification Authority (CA); and
- documenting a rough timeline for the C&A process.

Test and Evaluation Master Plan (TEMP)

The test and evaluation of IA requirements is an integral part of the overall Test and Evaluation (T&E) process. DoD Instruction 5000.2 directs that IA testing be conducted during both Developmental Test & Evaluation (DT&E) and Operational Test & Evaluation (OT&E). Key considerations for planning, coordinating, and executing IA testing include identifying sources of IA requirements, integrating C&A activities, and determining IA considerations for the TEMP. It is important that IA be adequately addressed in the TEMP to include IA roles and responsibilities, test strategies and summaries, and special resources. For example, the DAA should be identified and the Interim Approval to Operate (IATO) or Approval to Operate (ATO) should be included as entrance criteria for appropriate test events.

Acquisition strategy

To adequately address IA in the Acquisition Strategy document, the following should be included:

- **IA technical considerations—**Such as requirements governing Commercial-Off-The-Shelf (COTS) IA and IA-enabled products and Government-Off-The-Shelf (GOTS) IA or IA-enabled products
- **IA schedule considerations—**Such as IA C&A timeline and key milestones that are integrated into a program's TEMP
- **IA cost considerations**—Such as developing and procuring IA solutions, T&E, C&A of the IA architecture, and Operation & Maintenance (O&M) costs related to maintaining the system's security posture after deployment
- **IA funding considerations**—Including all IA lifecycle costs
- **IA staffing and support issues**—Ensuring that the program is adequately staffed to support IA requirements and has an appointed Information Assurance Manager (IAM) in accordance with DoDD 8500.1
- **As required**—Any other significant acquisition IA issues

These items are representative considerations of Acquisition Strategy IA and are provided solely as examples, but experience has shown that they are common to most programs. A PM should tailor and include text that addresses these items as appropriate.

IA assistance and supporting resources for the acquisition Program Management Office (PMO)

Defense-Wide Information Assurance Program (DIAP) Acquisition Team

The DIAP Acquisition Team is in place to assist and guide all stages of the acquisition process. It participates in acquisition program IIPTs, provides IA guidance to program IA Points of Contact (POCs), conducts early coordination reviews of Acquisition IA Strategies, and conducts formal DoD CIO reviews of Acquisition IA Strategies for CCA compliance. It also coordinates on program TEMPs, Acquisition Strategies, Acquisition Program Baselines, Acquisition Decision Memorandums, and JCIDS Documents and proposes and develops acquisition policy, guidance, and training related to IA.

Defense Acquisition University (DAU) learning module: IA for PMs

IA for PMs is an online DAU learning module that describes the importance of IA, the PM's responsibilities, and the steps for integrating IA into an acquisition program. It is available under the course name "Information Assurance" at the DAU Continuous Learning Center Web site at http://clc.dau.mil/kc/no_login/portal.asp

IA in acquisition section of the DAU acquisition community connection IT Community of Practice (CoP)

An IA in Acquisition section is available on the DAU Acquisition Community Connection IT CoP and contains the following sections:

- Introduction to IA in Acquisition
- IA in the Acquisition Life Cycle (The IA Roadmap)
- Emerging Issues (Coming Soon)
- Policy & Guidance Page
- Training Center
- Community Connection
- IA Resource Links
- What's New

The IA in Acquisition section of the DAU IT CoP can be found on the DAU Web site at http://acc/dau/mil. Once at the Web site, click on the "IT CoP" link.

Other resources

Other sources of IA assistance include the IA staff of a particular program's Program Executive Officer (PEO), major command or systems command CIO office, or Component CIO office. Also available is the IA Support Environment (IASE) Web site at http://iase.disa.mil, which contains DoD IA tools and resources, an IA document library, an "Ask the Experts" section, policy and guidance, a solutions database, and IA training.

Summary

Addressing IA in acquisitions may seem a daunting task, but the message is clear: plan and integrate IA in the acquisition process as early as possible and then follow through with best IA practices throughout the acquisition and beyond deployment. This is the most efficient and thorough way to promote the security of IT systems and support the successful achievement of program objectives. ■

About the Authors

Eustace King

Eustace King is the Deputy Director for Technology and Capabilities of the Defense-Wide Information Assurance Program (DIAP) within the Office of the Assistant Secretary of Defense for Networks and Information Integration (OASD/NII). His responsibilities include providing oversight for life cycle visibility of IA within DoD acquisition programs, ensuring that DoD maintains a viable IA research and technology development capability, implementing the Security Management Infrastructure to include the Public Key Infrastructure and Crypto Modernization, and ensuring the availability of technologies necessary for the rollout of the Global Information Grid (GIG). Mr. King is a 1975 graduate of Brooklyn College (City University of New York) and earned an MBA from Gongaza University in 1982. He is retired from the US Air Force after serving almost 25 years, is an NSA asset, and has been assigned to his current position for about four years.

Art King

Art King is a Managing Consultant with IBM Business Consulting Services and is currently supporting the Defense-Wide Information Assurance Program (DIAP). In this capacity, his primary role is to advance the integration of IA throughout the life cycle of defense systems acquisition. A member of the DIAP Acquisition Team since October 2002, he is responsible for engaging major acquisition programs and providing guidance regarding IA requirements and implementation. He also supports the development of IA related acquisition policy and guidance.

Mr. King is a former Electronic Warfare specialist and a retired Navy Supply Corps Officer. During his career with the US Navy, he served in operational, contracting, information management, and headquarters assignments. He is a member of the Navy Acquisition Corps and a qualified Joint Services Officer. Since his retirement, he has served as a consultant for Coopers & Lybrand L.L.P., PricewaterhouseCoopers L.L.P., and currently IBM Business Consulting Services.

Mr. King received his bachelor's degree from the University of the State of New York and a Masters degree in National Defense and Strategic Studies from the US Naval War College in Newport, RI.

Dominic Cussatt

Dominic Cussatt is a Senior Consultant with IBM, currently supporting the Department of Defense (DoD) Chief Information Officer's (CIO) Defense-Wide Information Assurance Program (DIAP). A member of the DIAP Acquisition Team since January 2004, he is responsible for engaging major acquisition programs to provide guidance regarding IA requirements and implementation. This includes reviewing program IA strategies and other key acquisition documentation for IA considerations and policy compliance. He also supports IA training development as a Subject Matter Expert (SME).

Mr. Cussatt has served as a consultant for Georgetown University, Litton, SAIC, and currently IBM Business Consulting Services as a member of its Security, Privacy, and Wireless Practice. He has provided consulting services to a broad range of government and private sector clients in the areas of IA, business and operations management, contract administration, organizational analysis, data systems analysis and design, and data center operations.

Mr. Cussatt received his bachelor's degree in International Politics with a Minor in Business Administration from the Pennsylvania State University.

Threats Posed by and to 802.11 Wireless Networks

by Michael Kershaw

Wireless networking, especially the Institute of Electrical & Electronics Engineers (IEEE) 802.11, also known as Wireless Fidelity (WiFi) networking, has become incredibly popular in the past several years. It has progressed from high-cost, relatively low-speed hardware in the late 90s, to a consumer commodity in 2001–02, to a nearly universal option built into consumer laptops and operating at more than 20 times its original speed. Unfortunately, despite many warnings, proof-of-concept attacks, and even mainstream media coverage, the risks presented by unauthorized or unplanned wireless networks and the risks posed to authorized networks are often poorly understood. This article addresses some of the most critical security and reliability flaws and discusses open-source software solutions for detecting unauthorized wireless networks and for auditing and monitoring authorized ones.

Unauthorized wireless networks

An unauthorized wireless network can completely expose a private network to sniffers and attackers. An unauthorized Access Point (AP) could be maliciously placed by an attacker who gained temporary physical access, but it is more likely to be placed by a worker wanting wireless access for their personal equipment (laptop, PDA, etc.).

The 802.11 protocol is designed to bridge seamlessly with wired Ethernet (802.3) networks: an AP dropped behind the firewalls will expose the entire network segment. If stronger methods of authentication—802.1x, per-port Media Access Control (MAC) filtering, or other link controls—are not used on the private, wired network, any user able to connect to the wireless is able to directly access any resource on the network, and the AP will replicate any traffic seen on the wired network (such as Windows file-sharing broadcasts), which can disclose yet other network resources.

Consumer-level 802.11 AP and router combination devices can also provide Network Address Translation (NAT) services for wireless clients. This means the AP appears on the local network as a single, non-bridging device while still giving wireless users full access to the private network. An AP in NAT mode is often far more insidious than a bridge, because network-monitoring tools will not alert that there are multiple clients on a single network port. Most equipment that can provide NAT and routing functionality can also arbitrarily set the Ethernet interface MAC address to clone an existing device on the network, such as a user's desktop.

If users have the necessary level of access to a workstation's hardware and operating system, an AP could also be placed on a second Ethernet card and operating-system-level NAT used to hide it entirely from network-side monitoring. This requires savvy users, but it can be a definite risk. This permits the system on the private network to respond to authentication queries and log-on methods but still exposes all private network services.

Detecting wireless on a network from the wired side is difficult or impossible, depending on the level of control exerted over the wired network. Scanning Content Address Memory (CAM) tables for known wireless MAC manufacturers' headers (the first three pairs of a MAC address indicate the manufacturer), preventing end-user access to administrator accounts and workstation hardware, strong port-level authentication, and MAC address restrictions can help reduce the risk, but the only sure way to detect wireless networks is with a wireless system, which will be discussed later in this article.

Secondary access to data or networks through unsecured default clients

Often overlooked, unsecured systems with wireless cards can present a threat to network and data security through secondary access routes. When a wireless client joins a network, it sends a "Probe Request" packet, which contains the name of the network it wishes to join—the Service Set IDentifier (SSID), the keyword "ANY," or a blank to indicate that it will join any public network. Typically, the drivers and utilities are configured to constantly search for networks and, if possible, to connect to them. This makes it easy on end users but very difficult for network administrators: any attacker within radio range can bring up an AP

and Dynamic Host Configuration Protocol (DHCP) server through which those clients can connect. Once a client has connected to the rogue AP, attacks can be launched against the client's operating system to gain access to data stored on the client or to access the private, wired network to which the client is connected.

There is no simple way for a network administrator to prevent clients from connecting to external networks so long as the administration of the hardware is in the hands of users (i.e., personal laptops, PDAs). Forcing users to routinely run updates and scan for malicious software can minimize windows of opportunity should an attacker provide a fake AP, as can educating users against leaving wireless cards active, but neither is a perfect solution.

Unauthorized access through official APs

Incorrectly configured official APs can create the same security holes in your private network as rogue, unofficial APs. The following are key factors to remember when adding wireless to an existing network:

- Always treat wireless networks as hostile, external network segments. An AP should never be placed directly on a secure network segment. Always place APs outside firewalls and require the same authentication procedures as would be required to access the network from a remote location. These authentication procedures should use strong methods, such as mandatory Virtual Private Network (VPN) connections to enter the private network.

- Most vendor security methods are not added security. When the initial wave of media coverage of 802.11 vulnerabilities occurred, many vendors implemented "security" measures that provided no significant benefit. The first of these was "weak Initialization Vector (IV) avoidance," which altered the Wired Equivalent Privacy (WEP) data generator to avoid generating keys that fell within the Fluhrer-Mantin-Shamir (FMS) Weak Key classification (to be discussed later). Unfortunately, the RC4 [1] stream-cipher algorithm has several other significant flaws, and this is by no means adequate protection from them. Other security measures attempt to modify the 802.11 protocol to hide the name of an AP. Called "SSID Hiding" or "SSID Cloaking," an AP no longer sends the network name in every beacon packet. In theory, this requires clients to know the network name before they can connect. However, the SSID field was never meant as a security mechanism, and the network name is returned in plaintext in the response from the AP when a client connection is accepted. By waiting for a legitimate client to connect, an attacker can instantly discover the name of the network and connect; a determined attacker can cause all clients on a network to disconnect and reconnect, forcing the SSID disclosure. Some vendors also implement MAC address filtering in an attempt to keep authorized clients from connecting to the AP at all. Again, an attacker has only to clone the MAC address of an authorized client, and this measure is bypassed. The MAC addresses of clients are passed in the clear, even if network-level encryption is used. Combined, these methods can provide some minimal protection against casual intruders, but they should never be mistaken for real security.

Active attacks against 802.11 networks

802.11 wireless networks are vulnerable to a variety of attacks that range from hijacking sessions on unsecured networks to Denial-of-Service (DoS) attacks that will cripple any network no matter how secure. 802.11 cannot be made into a medium hardened against DoS attacks and should never be trusted for any critical task. The critical flaws in 802.11 networking fall into two main categories:

- Flaws afflicting any radio network, including passive sniffing and interference from external sources

- Flaws in the protocols, implementation, and built-in encryption of 802.11 networks that expose all networks to protocol-level interference, spoofing, and man-in-the-middle attacks at the session layer

The basic nature of radio permits both undetectable snooping and remote interference with the operation of a network. With proper equipment and a clear line of sight, network data can be gathered from literally miles away. Wireless networks are also vulnerable to interference and noise, malicious and not. IEEE 802.11b and 802.11g networks run on the

2.4-GHz unlicensed spectrums, which are shared by a plethora of consumer electronics devices (e.g., microwave ovens, baby monitors, and some cordless phones), and the 802.11a 5-GHz band is shared with newer consumer electronic devices attempting to escape the congestion of the 2.4-GHz band.

The 802.11 protocol also lends itself to simple attack methods. While there is some protection afforded to data frames (link-layer WEP encryption and minimal checksumming to prevent tampering), no such protection is given to management frames, which define and control the network. All management authentication is based on a MAC address, and all management data is sent unencrypted. All 802.11 networks, no matter how secure the data layer is, are vulnerable to management-level attacks:

- **Client hijacking**—The 802.11 protocol was designed to permit clients to seamlessly roam between physical APs by using the SSID to define a single, continuous network. If a client is no longer able to connect to the AP it was using, it will begin scanning the available channels for another AP. By creating an AP with the same name as that of the legitimate network and by forcing a client to disconnect from a legitimate AP (either by causing radio interference or by spoofing the AP and sending a disconnect packet), an attacker can hijack a client and relay all its data to the legitimate network, all without the client ever being notified that it has changed networks. This attack can be mitigated by using link-level authentication, such as 802.1x, and strong network encryption, such as a VPN. It is crucial that the VPN software pre-share the network authentication keys. Proof-of-concept code has existed for two years; it can perform an automatic man-in-the-middle attack against VPN software that does not have pre-validated keys (AirJack, no longer available from the developers site), completely negating the encryption.

- **Client disassociate and de-authenticate DoS**—If an attacker spoofs the legitimate AP and continually sends de-authentication and disassociation frames, clients will continually be booted from the network before they're able to usefully exchange data. This DoS attack can cripple any network within transmission range of the attacker for as long as the attacker wishes to continue the attack. There is no way to mitigate an attack of this type short of finding the attacker physically and halting it.

- **AP disassociate and de-authenticate denial of server**—By spoofing a single client, an attacker can continually request that the AP sever the connection. This is the same mechanism as that discussed in the previous attack, but it targets only a single client.

- **Power-save exploits**—If an attacker spoofs a client that is in power-save mode and requests all packets be held pending for that client, the packets will be delivered while the client isn't watching for them. Similarly, an attacker can spoof a client and notify the AP that it will be entering power-save mode. This often causes packets at the AP to queue up and otherwise generally disrupts traffic.

- **Firmware exploits**—Firmware, as with any software, can contain bugs, which are sometimes critical. Specifically, most firmware releases for Prism2 cards (an 802.11b chipset, licensed and re-labeled by many different manufacturers) and Orinoco (owned by several companies and licensed for other products, such as the original Apple Airport) contain vulnerabilities that cause a crash of the card on receiving a poisoned frame. An exploited card will freeze, often causing the host operating system to also freeze and fail. Other manufacturers may also be plagued with similar hidden problems.

Monitoring wireless networks

Sniffing 802.11 networks is different from sniffing normal Ethernet. Under normal operation, the 802.11 network layer is hidden from the operating system. The driver only passes the contents of data frames to the operating system, formatted as 802.3 Ethernet. Normally, a packet sniffer places a card in promiscuous mode, whereby the card no longer filters packets that are not destined to its MAC address. Placing an 802.11 card in promiscuous mode will typically have no effect or will return all data frames from the currently associated network but will not disclose data from other networks or the 802.11 layer itself.

To sniff the 802.11 layer, both your card and drivers must support a mode of operation called "rf monitor" or "rfmon." While in monitor mode, a card cannot be part of a network or transmit, but it will report all packets on the channel, including 802.11 management frames, data frames from any network, and encrypted data frames. Currently, the best platform for monitoring wireless is Linux, as it has the greatest selection of drivers that support monitor mode. Some chipsets also work on Berkeley Software Distribution (BSD) systems, but monitor mode is currently not as pervasive or as well supported. There are no current, free, or legally unencumbered drivers for Windows, and it is not a good choice for monitoring wireless with open products.

Some programs, such as NetStumbler (http://www.netstumbler.com), detect wireless networks through a different method: firmware scanning. The firmware of wireless cards builds a list of available networks that have responded to probe requests or, in some cases, that the card has seen beacons from. Scanning mode will disclose public networks in an area but cannot detect hidden networks or networks that are out of transmission range of the card that is scanning. Also, scanning mode cannot return data from networks, as it only queries the firmware for networks that have announced themselves. The core toolkit for discovering and analyzing wireless networks is three open-source programs:

- **Kismet** (http://www.kismetwireless.net)—Kismet is an 802.11 sniffer, a layer-2 Intrusion-Detection System (IDS), and network tracker. It can be used as a mobile sniffer in combination with a Global Positioning System (GPS) and includes mapping software to plot wireless networks and signal levels. It can also be used as a stationary IDS to monitor for unauthorized networks or attacks against an AP. Kismet will homogenize the packet streams from different drivers into a standard IEEE 802.11 packet file.

- **TCPDump** (http://www.tcpdump.org)—TCP Dump is a standard packet sniffer for a variety of network types. TCPDump is excellent for testing connectivity and quickly reviewing the packets going over a network.

- **Ethereal** (http://www.ethereal.com)—Ethereal is a graphical packet sniffer with decoders for hundreds of protocols and link types. Ethereal is best for dissecting a previously captured file, monitoring application protocols, tracking Transmission Control Protocol (TCP) connections, and other data analysis.

Kismet

Kismet is designed to capture data from various 802.11 cards, homogenize the various packet headers and capture types into a consistent IEEE 802.11 packet stream, and automatically sort the data into networks. Network, client, and GPS data can be exported to eXtensible Markup Language (XML) for translation to other reports, and a live packet stream can be exported through a POSIX*-named pipe.

Kismet is a completely passive network sniffer; with the exception of cards with buggy firmware (noted in the readme), it will not transmit frames while sniffing, and its operation is not normally detectable. It can also automatically disclose the SSID of hidden networks by monitoring for responses from the network when a client joins, decrypt WEP encrypted packets in real time if the key is known, and alert on access-point configurations that match known factory defaults. If the network data is unencrypted or if a known key exists for the network, Kismet will also attempt to guess the IP range used on the network by analyzing data frames. The Kismet IDS system comprises two methods of detecting problems:

- ***Fingerprinting*** uses known packet data that indicates there is an attack or other problem. Unfortunately, not all attacks are as easily identified; many DoS attacks involve packets that normally are perfectly legitimate in most circumstances.

- ***Trend-based alerts*** trigger on the frequency or order of packets that create an attack. Kismet can currently detect attacks against the firmware of some brands, attempts at network spoofing, DoS floods, specific NetStumbler versions, and generic stumbler activity.

Kismet consists of four core components:

- ***Packet-capture engine and network server (kismet_server)***—The capture engine performs the actual packet gathering (from standard lipcap-based interfaces or more exotic methods, such as the WSP100 remote sensor, Kismet remote-capture drones, or custom drivers), packet filtering, packet homogenization, and logging and provides the TCP server to which thin clients connect. The server can be run headless to log packets and network events, or it can support any number of clients to remotely report data.

- ***Default front-end client (kismet_client)***—The default client Kismet provides is a Ncurses-based display that will work on any text terminal. Other third-party clients provide full graphical views of data sent by the server, and Perl modules, available on the Kismet site, make writing custom clients for logging or forwarding event data trivial.

- ***Minimal capture engine that exports data captured from wireless cards over a wired network connection***—This lightweight engine can run with extremely minimal hardware requirements, including embedded devices, such as the Linksys WRT54G AP and other low-power, embedded devices. The drone architecture is designed to permit cheap, stationary sniffing hardware to be placed throughout a building to report all packets to a central collection point for analysis and decoding.

- ***Mapping component (GPSmap)***—This component processes the gps-xml and network-xml files generated by the Kismet server. GPSMap processes samples, screens garbage points, calculates best-guess network centers, and plots the output on maps downloaded from publicly available map sources such as MapBlast, the US Census Tiger database, MapPoint, Terraserver, and others. GPSMap can plot convex network hulls of all seen sample points, interpolated calculation of probable network coverage, network-center guessing, plotting of every sample point, guessed network ranges, and coloring by channel and encryption status.

Crucial Kismet configuration

Kismet uses the standard GNU auto-configuration scripts (./configure) to configure the source, and has a typical “make” and “make install” process. On single-user systems (such as laptops), Kismet may be installed suid-root by using “make suidinstall,” but it should never be installed suid-root on a multiuser system. The Kismet configuration file has many options; however, the following are most critical:

- ***suiduser***—Once Kismet has bound to the capture devices, it will leave a minimal root process running to manage the channel hopping and drop privileges to a user that are specified to protect against any potential exploits.

- ***source***—The source lines define what capture sources Kismet listens to. A complete list of supported capture sources can be found in the readme file. If multiple source lines are specified, Kismet will listen to them all and multiplex the data into a single packet stream.

- ***alert***—Alert configuration lines control the intrusion detection alerts and the rate at which multiple alerts occur. Alert rates can be limited by type, rate per timeframe, and burst rate: For example, an alert for possible spoofed APs could be limited to five alerts per minute with a maximum of three alerts in one second.

- ***logtypes***—The logtypes line controls the log files generated by the Kismet server. Log files can be dump (pcap-formatted packet dump file), network (plaintext list of networks and information), CSV (Comma Separated Values' table of network information), XML (network-xml file of the detected networks), weak (pcap-formatted file of packets that meet the FMS weak WEP criteria, suitable for loading into airsnort), Cisco (plaintext file of Cisco discovery information), and GPS (XML file of GPS coordinates for plotting).

- ***logtemplate***—The logtemplate format string controls the naming of log files. By default the files are written to the current directory when Kismet is started, but by prefixing the logtemplate configuration with a path, logfiles will always be put in a specific directory. When editing the logtemplates, make sure the suiduser has permissions to write to the target directory.

Kismet in action

Starting Kismet will place the capture sources and begin capturing packets.

Figure 1 shows Kismet in normal network display mode. From this screen, you can receive detailed network information, client lists, summaries, alerts, and all other relevant data. In this configuration, green networks are encrypted, yellow networks are unencrypted, red networks match factory-default configurations, and blue networks are networks with hidden SSIDs that have been decloaked.

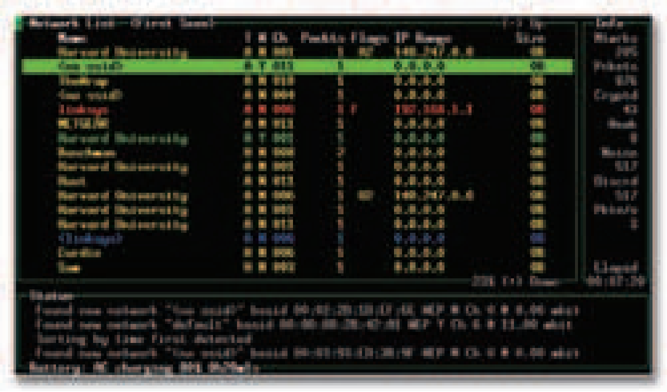

Figure 1. Normal mode Kismet

Figure 2 shows overall statistics for the current sniffing run, including packet rates, network statistics, and a histogram of channel usage.

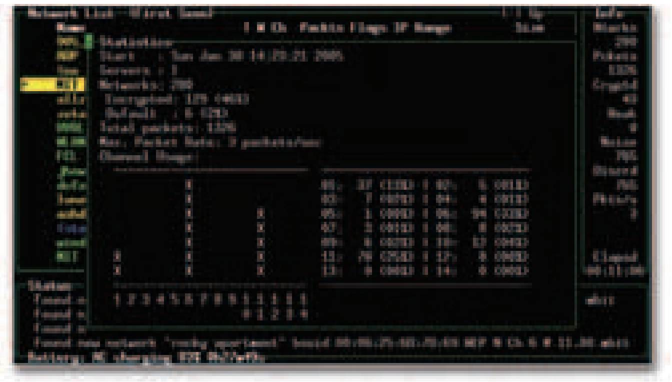

Figure 2. Overall statistics

Figure 3 shows how Kismet tracks extensive information for each seen network. This data is also logged into the network files and the network-xml file. Manufacturers are matched by the first three octets of the MAC address, and attempts are made to match the exact model by using more digits of the MAC address.

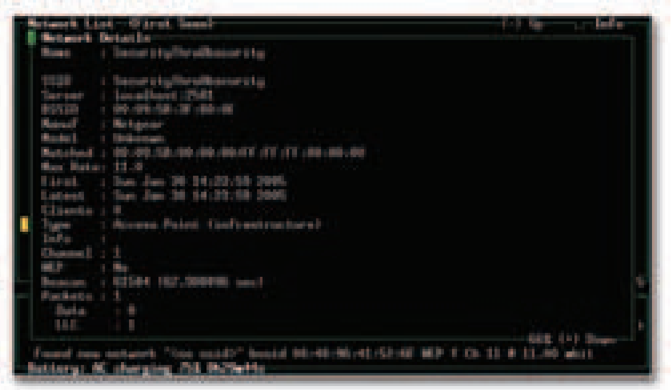

Figure 3. Network details

Figure 4 shows how Kismet can detect wired clients that are broadcast on the wireless network, wireless clients joining the network without sending data, and established clients talking on the network. Currently active clients are marked with an exclamation point.

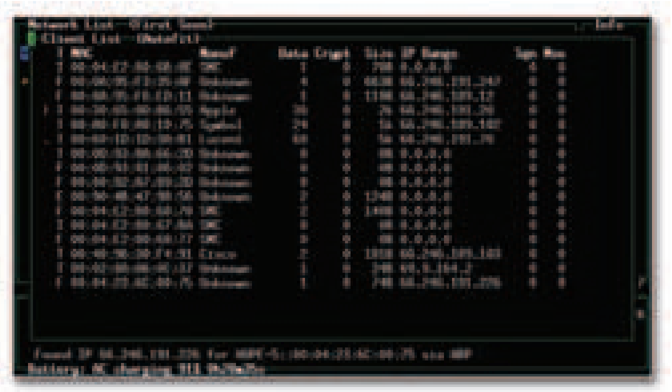

Figure 4. Client lists

Figure 5 shows more detailed information about the packets in view. This is by no means intended to replace full-packet analysis tools such as Ethereal but can be valuable in analyzing a problem.

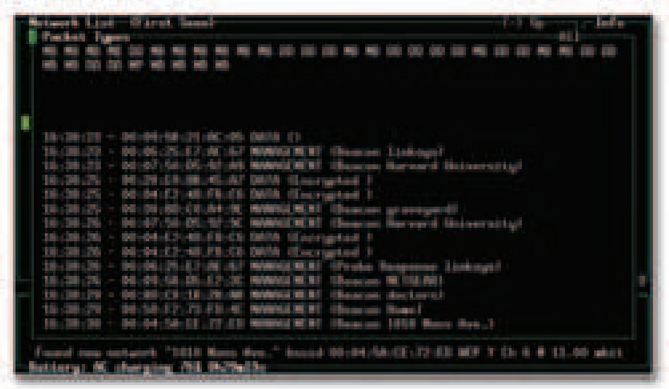

Figure 5. Packet information

Figure 6 shows how alerts are displayed in the status box and in a dedicated alert window. The alert's time, classification, and details are reported. In this case, the alert system was configured to report trends that indicate scanning clients, such as NetStumbler. ■

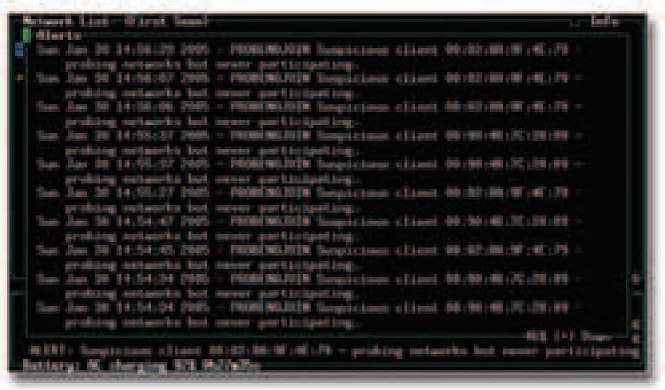

Figure 6. Alert

References

[1] Ron's Code 4, a Rivest, Shamir, Adleman (RSA) Variable-Key-Size Encryption Algorithm by Ron Rivest

About the Author

Michael Kershaw

Michael Kershaw has been involved with wireless and wireless security for four years, and is the author of the open source wireless sniffer and intrusion detection program Kismet (http://www.kismetwireless.net) as well as random additional wireless projects and deployments.

letters to the director

Recently, the IATAC Technical Inquiry Specialist received the following question, which I want to share with our readers:

I have been hearing about something called a Research and Engineering Portal. Could you please tell me what you know about it?

The Research and Engineering (R&E) Portal is the focal point for obtaining information on current and historical research and engineering activities within the Department of Defense (DoD), and is available to DoD employees and their contractors. The Portal is sponsored by the office of the Director of Defense Research & Engineering (DDR&E) and maintained by the Defense Technical Information Center (DTIC). The R&D Portal includes such information as data from systems that focus on the areas of Financial Management, Strategic Planning, and Congressional Reporting as well as information on areas of strategic importance and current initiatives within DDR&E.

Additionally, within the Portal, you will find tools to facilitate collaboration, communication, and reuse of information and artifacts, with the added benefit of robust text searching tools to query the wealth of DoD research and engineering information held by DTIC and other services and agencies.

For additional information about the R&E Portal and to learn how to register, please visit http://www.dtic.mil/ or contact either rdte_help@dtic.mil or iatac@dtic.mil. ■

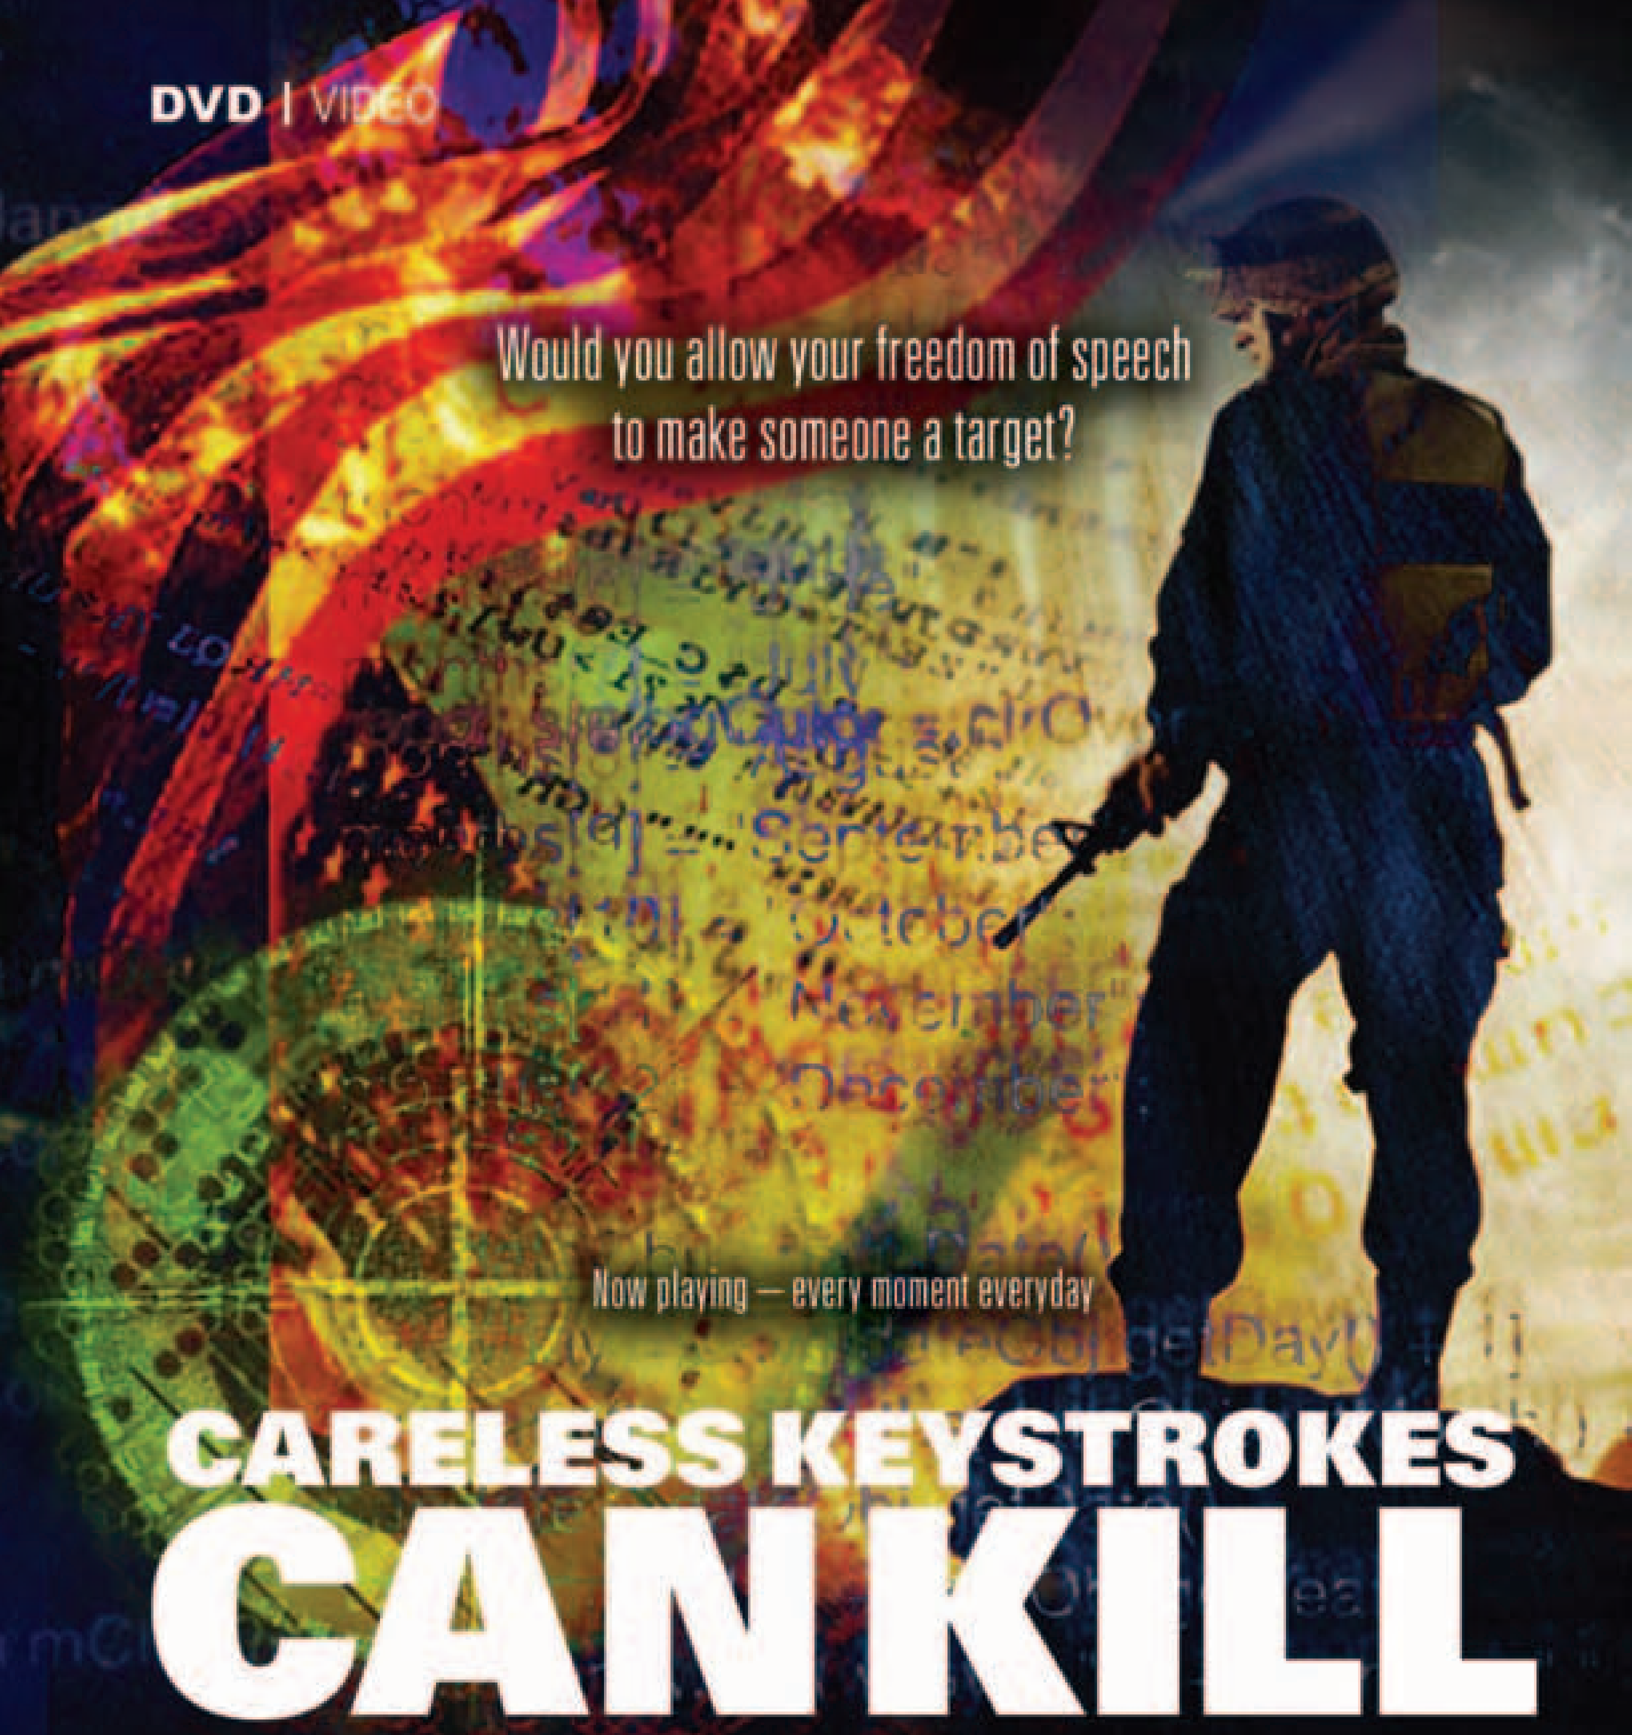

DVD | VIDEO

Would you allow your freedom of speech to make someone a target?

Now playing — every moment everyday

CARELESS KEYSTROKES CAN KILL

Learn how to see your surroundings through the adversary's eyes. . . to avoid the everyday mistakes that can add up to tragedy. Don't become the adversary's best friend.

Would You Allow Your Freedom of Speech to Make Someone A Target?

That's the question raised by *Careless Keystrokes Can Kill*, a video produced by the U.S. Strategic Command (USSTRATCOM) and the Joint Task Force for Global Network Operations (JTF-GNO) which premiered July 7 at the Component Commanders Conference in McLean, Virginia. Viewing this video could literally save your life.

Based on a tagline coined by recently retired Lt Gen Harry D. Raduege, Jr., USAF, the video focuses on the importance of operations security (OPSEC) for the Global Information Grid (GIG). The 15-minute piece features a series of vignettes that traces the origins of a believable military tragedy, explains how the tragedy could have been averted, and demonstrates ways to prevent mistakes that can unknowingly provide sensitive, unclassified information to the adversary. A personal appeal from Gen James E. Cartwright, USMC (Commander, USSTRATCOM) is included to underscore the video's message.

Careless Keystrokes Can Kill spotlights the all-too-human mistakes and slips made by bloggers, warfighters using 3rd party e-mail while deployed, and users of wireless handheld devices and cell phones. The importance of physical security is also covered. In addition, the video addresses the frightening statistic that an average of eighty percent of the information needed to sabotage a mission can be obtained through unprotected, open sources.

This video is intended for viewing by warfighters at every level from enlisted to command, Department of Defense personnel, and civilian contractors involved in information assurance activities.

It is hoped that the slogan *Careless Keystrokes Can Kill* will become for our technologically oriented time what *Loose Lips Sink Ships* was for the World War II generation.

Careless Keystrokes Can Kill has been designated as FOR OFFICIAL USE ONLY, Distribution Statement D. To inquire about copies of the video, please visit the Interagency OPSEC Support Staff—IOSS—at www.ioss.gov or the Joint Information Operations Center at www.jioc.smil.mil).

IATAC Conference and Event Planning

Experienced Assistance for Your Classified or Unclassified Event

Are you a government client in need of planning and hosting assistance for an upcoming conference? Look no further…

IATAC's conference and event planners provide the assistance you need.

Since 1998, we have offered a full range of services to support classified and unclassified conferences, meetings, and other gatherings for groups ranging from 20 to 300+ participants. From site selection and registration to catering and security requirements coordination, we can plan and execute an event that complies with government conference regulations and provides a high level of customer satisfaction. All members of our staff hold active security clearances ranging from Secret to Top Secret/SCI.

Services are available to all government clients regardless of whether or not they are currently affiliated with the IATAC contract. Support can be arranged through Technical Area Tasks (TATs) subscription accounts with payments via Military Interdepartmental Purchase Requests (MIPRs), if applicable.

Our experienced planners offer service and support for all phases of your event.

Before the event

- Site selection
- Budget oversight
- Contract negotiation
- Secure online registration and payment
- Graphics support
- Audio/visual coordination
- Agenda development
- Sponsorship/exhibitor solicitation
- Marketing and promotion
- Security requirements coordination (classified events)

During the event

- Check-in and registration
- Collection of registration fees
- Note-taking (session minutes)
- Speaker assistance
- Problem resolution
- Catering coordination

After the event

- After-action report
- Conference surveys and evaluations
- Distribution of conference proceedings
- Reconciliation of invoices and registration fees

Want more information?

To find out more about IATAC's conference and event planners and what they can do for you, please contact:

April Perera
Director, Conference and Event Planning
703/289-5699

Avery-Lynn Dickey
Conference and Event Planner
703/289-5559

Team e-mail: iatac@dtic.mil

Examples of recent events

GO/FO/SES Global NetOps Conference,
July 2005

Fourth Intel Support to CND Conference,
March 2005

JTF–GNO Reporting Working Group,
February 2005

Joint Task Force for Global Network Operations (JTF-GNO) Component Commanders Conference,
January 2005

DoD Defense Continuity Conference,
September 2004

Treasury IT Security Conference 2004: Making the Grade,
June 2004

Federal PKI Deployment Workshop 2: Federal Credentialing and Beyond,
May 2004

The Political/Military Dimensions of Cyber Security,
March 2004

Second Intel Support to CND Conference,
February 2004

Intel Support to CND Conference,
August 2003

Federal PKI Deployment Workshop,
March 2003

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	□ No. 4
Volumes 8	□ No. 1	□ No. 2		

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

September

2005 Information Assurance Conference & Exposition

September 7–8, 2005
Sheraton Premier Hotel at Tysons Corner, Vienna, VA
https://www.technologyforums.com/amc/index.asp

Air and Space Conference and Technology Exposition

September 12–14, 2005
Marriott Wardman Park Hotel, Washington, DC
http://www.afa.org

Biometric Consortium Conference 2005 (BC2005)

September 19–21, 2005
Hyatt Regency Crystal City, Arlington, VA
http://www.nist.gov/public_affairs/conf-page/050919.htm

October

Strategic Space 2005

October 4–6, 2005
Qwest Convention Center, Omaha NE
http://www.stratspace.org/information/index.cfm/

2005 Homeland Security Summit

October 12–14, 2005
Holiday Inn Rosslyn at Key Bridge, Arlington, VA
http://www.homelandsecurityweb.org/HSS/

MILCOM 2005

October 17–20, 2005
Atlantic City Convention Center, Atlantic City, NJ
http://www.milcom.org/2005/

TechNet Europe 2005

October 17–20, 2005
Atlantic City Convention Center, Atlantic City, NJ
http://www.afceaeurope.org/html/technet_europe.html

InfoTech 2005

October 18–20, 2005
Dayton Convention Center, Dayton, OH
http://www.afcea-infotech.org/

2005 Homeland Defense Symposium

October 24–27, 2005
Broadmoor Hotel, Colorado Springs, CO
http://www.hldsymposium.org

November

Physical Security: Securing America One Building at a Time

November 2-3, 2005
NRECA Executive Conference Center, Arlington, VA
http://www.homelanddefensejournal.com/conf_criticalinfrastructure.htm

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