13/2

Volume 13 Number 2 • **Spring 2010**

IAnewsletter

The Newsletter for Information Assurance Technology Professionals

Cloud Computing:
Silver Lining or Storm Ahead?

also inside

IATAC

contents

About IATAC and the *IAnewsletter*
The *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), and Director, Defense Research and Engineering (DDR&E).

Contents of the *IAnewsletter* are not necessarily the official views of or endorsed by the US Government, DoD, DTIC, or DDR&E. The mention of commercial products does not imply endorsement by DoD or DDR&E.

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

IATAC Director: Gene Tyler
Inquiry Services: Peggy O'Connor

If you are interested in contacting an author directly, please e-mail us at Iatac@dtic.mil.

***IAnewsletter* Staff**
Art Director: Tammy Black
Copy Editor: Kali Wilson
Designers: Michelle Deprenger, Dustin Hurt
Editorial Board: Dr. Ronald Ritchey, Angela Orebaugh, Gene Tyler, Kristin Evans, Al Arnold

***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 email address (soft-copy receipt), please contact us at—

IATAC
Attn: Peggy O'Connor
13200 Woodland Park Road
Suite 6031
Herndon, VA 20171

Phone: 703/984-0775
Fax: 703/984-0773

Email: *iatac@dtic.mil*
URL: *http://iac.dtic.mil/iatac*

Deadlines for Future Issues
Summer 2010 May 8, 2010

Cover design: Tammy Black
Newsletter design: Donald Rowe

in every issue

IATAC Chat

Gene Tyler, IATAC Director

In early February, I had the opportunity to attend the Information Assurance Symposium (IAS) in Nashville, TN. I always look forward to attending this event because it brings together folks who truly care about information assurance (IA). I am always excited to converse with colleagues interested in solving tough IA problems ahead, and yet again, the IAS did not fail; I enjoyed talking with people about some of the newest innovations currently changing our field.

One topic that seemed to dominate the conversations I had with various colleagues and subject matter experts at IAS was cloud computing, and as this edition of the *IAnewsletter* reflects, this topic is getting a lot of well-deserved attention, for a multitude of different reasons. Cloud computing is revolutionizing how organizations are constructing their networks and systems; it is changing how organizations invest in their information technology infrastructure; and it is forcing organizations to reconsider how they secure critical information—security is critical and at the forefront of cloud computing

But what, exactly, is cloud computing; and how do you ensure information security in the cloud computing environment? Dr. Bret Michael and Dr. George Dinolt, of the Naval Postgraduate School (NPS), address some of these questions in their article, "Establishing Trust in Cloud Computing." They argue that a lot of discovery is necessary before the IA community can fully understand cloud computing, its benefits, and more importantly, its weaknesses. I believe they say it best in their statement, "It is unclear whether the current set of [cloud computing] services is sufficiently secure and reliable for use in sensitive government environments." They advocate a cautious approach to implementing cloud computing capabilities across the government and, in particular, the Department of Defense (DoD). However, these subject matter experts remain optimistic, which is why they are excited about the research and investigation NPS is doing to identify methods of securing cloud-based systems.

On the other hand, some organizations are beginning to successfully implement cloud computing already. Most notably, the Defense Information Systems Agency (DISA) successfully developed the Rapid Access Computing Environment (RACE), which is a cloud-based system. Not only has DISA successfully implemented RACE, but, as the authors point out, "certification and accreditation policy has been adapted to allow organizations to use RACE cloud resources, thereby quickly connecting to the cloud while complying with DoD requirements." Munjeet Singh and Troy Giefer remain deeply involved with DISA as it implements cloud solutions, and as a result, their article, "DISA RACE: Certification and Accreditation for the Cloud," provides a different perspective on cloud computing and its advantages.

As these two articles suggest, there is a lot of debate over cloud computing, the advantages it offers, and the risks it presents. I hope the articles presented in this edition of the *IAnewsletter* also provide you with various perspectives on cloud computing so that you feel inspired to enter into the dialogue. I ask you, is cloud computing the silver lining to computing, and should we storm ahead in implementing it across various organizations? Or might it weaken our computer network defenses and result in a potential storm of malicious attacks in the future?

In addition to cloud computing, I invite you to look at the various other articles in this edition that highlight the following topics, also discussed at IAS: insider threat; Web 2.0 Security; social media and its use in DoD; vulnerability assessments; defending the Global Information Grid; and our industry expert contributes a very interesting article on public/private partnerships. As I always remind our readers, we are interested in your perspectives and welcome your contributions to this publication. We know our readers are the very subject matter experts who are analyzing and experimenting with innovative solutions like cloud computing. Feel free to contact us at iatac@dtic.mil with your perspective on the cloud debate!

Establishing Trust in Cloud Computing

by Dr. Bret Michael and Dr. George Dinolt

In the aptly titled article, "Cloud Assurance Still Missing," Allan Carey wrote, "The security problems that organizations face related to cloud computing are the same as those related to virtualization—but even more so." [1] He goes on to say, "Information assurance practitioners already have most of what is needed to make an informed set of decisions about cloud computing." [2] We would argue that the security problems go well beyond the use of virtualization in distributed systems. In this article, we discuss the need for asking critical questions about the security implications of cloud computing. Answers to our questions are not readily apparent, even though viewing computing as a utility, similar to that of providing water or electricity on a for-fee basis, dates back to at least the 1960s. [3]

As we pointed out in a recent article, [4] what has changed over time is the advancement of the underlying technology, including cheap, fast central processing units (CPUs), low-cost random access memory (RAM), inexpensive storage, and the high-bandwidth standardized communication needed to efficiently move data from one point to another. Additionally, considerations, such as the economies of scale involved in building very large data centers, nudged organizations to consider cloud computing as a vehicle for maintaining their competitive edge.

A recent technical report published by the University of California, Berkeley, states that there is no commonly agreed upon definition of cloud computing. [5] Instead, a definition is emerging as the various organizations that are developing cloud services evolve their offerings. In addition, there are many shades of cloud computing, each of which can be mapped into a multidimensional space with the dimensions being characteristics, service models, and deployment models. [6]

Cloud computing is a metaphor for giving Internet users a growing collection of computer system resources and associated software architectures to provide application services. [7] The applications include processing and application integration, storage, and communications services. Cloud services are typically available on demand and are charged on a usage basis. Often, what the user sees is an application instead of a particular computer. The services are commonly described as:

- **PaaS (Platform as a Service)**—the cloud provides hardware resources, typically virtual machines, which can be loaded with the users, operating system and software;
- **IaaS (Infrastructure as a Service)**—the cloud provides an infrastructure including (virtual) platforms, networking, *etc.* on which applications can be placed;
- **SaaS (Software as a Service)**—the cloud provides software applications.

Amazon's Elastic Compute Cloud (EC2) is an example of these services. [8] Google also provides enterprise-level integrated application services such as email, appointment calendars, text processing and spreadsheets. [9]

The claimed advantages for an enterprise are that it does not require an investment in computer resources, infrastructure, administration, *etc.*: the purveyor of the cloud provides these resources. The user or enterprise only pays for the resources "consumed." In the Department of Defense (DoD), we have seen the introduction of infrastructure services on demand provided by the Defense Information Systems Agency's Rapid Access Computing Environment (DISA RACE). [10] Where available, the cost of developing and maintaining specialized applications can be shared among the users of that application. In theory, there is an advantage in having large-scale resources shared among a large class of users. However, this has yet to be borne out. [11] There are, of course, applications that require a large number of resources. Google Search is one such

example. It appears that Google, Amazon, and others are attempting to leverage their ability to construct such a system into other environments.

We can argue that it is not a matter of whether cloud computing will become ubiquitous but rather what we can do to improve our ability to provide cloud computing users with assurance that the cloud services and infrastructure provide appropriate security functionality. Cloud computing providers should supply their customers with an appropriate level of security transparency to alleviate customers' reservations about the security and privacy afforded by the cloud. [12] How much transparency is enough? How do we provide for transparency of cloud resources (i.e. determining the cloud in which customer data resides)? Is there a tipping point at which additional levels of transparency would only serve to help malefactors compromise services and datacenters?

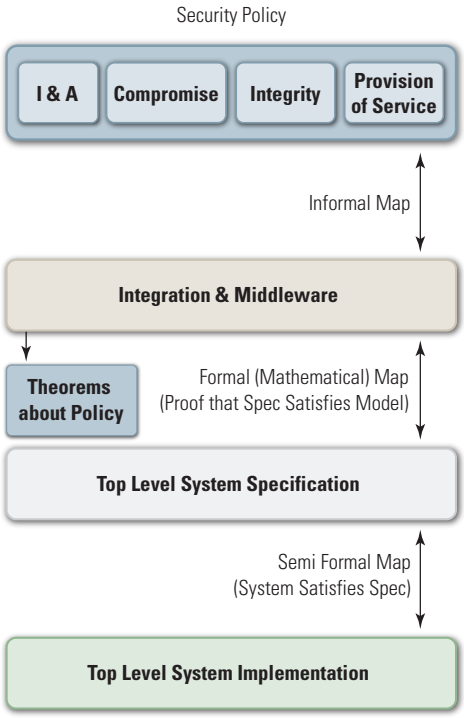

Figure 1 Process for Integrating Security Into the Cloud

In addition, as users and developers find new ways of applying cloud technologies, there will be new expectations about security and privacy. For instance, Twisted Pair Solutions of Seattle proposes to provide cloud computing resources for state and local agencies to link up disparate public safety radio systems (*e.g.*, police, fire, or ambulances)—a novel but difficult-to-predict usage of cloud computing, but also a usage that makes the cloud part of mission- and safety-critical systems. [13] The expectations for security, privacy, reliability, and quality of service and so on will be different in some respects for Voice over Internet Protocol (VoIP) radio systems than for the cloud's social networking aspects. This raises the question: how do we manage risk when we do not fully understand what we are trying to protect or guard against?

The fluid nature of cloud computing makes it a moving target, even when trying to determine the questions we should be asking to improve the security and privacy clouds afford. However, we can ask fundamental questions like: are the current architectures adequate for building trusted clouds? If not, what types of software system architectures do we need? Consider, for instance, the possibility that an organization might opt to fully outsource its computing infrastructure and data center to the cloud, retaining only thin clients within the organization. How do we make the thin client user terminals and the communications infrastructure secure?

DoD Enterprise Computing

What is our motivation for jumping feet first into asking hard questions about cloud computing? The growing importance of cloud computing makes it increasingly imperative that security, privacy, reliability, and safety communities grapple with the meaning of trust in the cloud and how the customer, provider, and society in general gain that trust. Consider the initiative of the DoD Enterprise Services & Integration Directorate to make the DoD Storefront Project a reality. The Storefront consists of a cloud-based set of core and specialized applications that users can discover through an application marketplace and which share an identity management framework. How will DoD provide security for the Storefront? It is more than a matter of having an identity

management framework. The obvious security concerns include data integrity, data availability, protection of personally identifiable information, data protection, data destruction, and communications security.

Moving beyond the Storefront concept, as the federal government migrates its data and applications to the cloud, issues regarding cross-domain resource sharing will arise within the cloud. For instance, how will DoD link its clouds to those of other agencies? Will a DoD user, authenticated to enter the DoD cloudsphere, be trusted to access services owned by the Department of Homeland Security (DHS)? Is there a need for a federal-wide cloud infrastructure and common set of security services? How will data be shared among the various different types of cloud?

Information Assurance

At the Naval Postgraduate School, a major thrust of our research on cloud computing is to investigate the security policies, models, and appropriate architectures to provide security for entities/users of cloud computing resources. Although cloud computing may appear to provide reasonably well understood operating system and application resources, cloud resources are distributed in space, time, and scale in ways that were never envisioned in the operating-system world. The current architectural approaches, especially those concerning security, may not scale to the much larger cloud computing approaches. In addition, the approaches for assuring operating system security functionality are not necessarily appropriate. It is unclear whether the current set of services is sufficiently secure and reliable for use in sensitive government environments. Current security claims are somewhat limited.

One of the fundamental problems with adopting cloud computing is providing not only security resources but also assurances that those resources are correctly implemented and maintained within the cloud. Several vendors have formed the Cloud Security Alliance (CSA). [14] In the report titled *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1,* CSA provides its take on some of the security issues related to cloud computing. [15]

In the report, security properties are described as essentially the same set of properties that a user expects to see with a self-hosted system. These include the usual:

- Identification/Authentication
- Privacy
- Integrity
- Provision of Service.

They view assurance as an audit of the function's implementation, that is, the cloud systems' administrators and implementers have used 'best practices'. Other than the notion that encryption is used to protect the data, there is little information that defines 'best practices.' There is, however, some form of key management included that provides potentially strong identification/authentication, as well as some form of data integrity/recovery facility. The security architecture proposed is essentially a layered operating system application. It consists of a network layer interposed between application programming interfaces (APIs) and the underlying operating system infrastructures. 'Trusted computing' is only mentioned at the hardware/operating system level. Additionally, the CSA paper enumerates several security issues that should be addressed by the cloud-style service provider, but does not provide any insight on security policies/models, interfaces or potential solutions.

To provide an example of some of the potential issues, Google supports "Google Apps." [16] Google Apps applies the usual discretionary access controls to the resources it provides – files, calendars, address lists, *etc.* To make life easier, Google provides tools that integrate their identification and authentication systems into the enterprise providing single sign-on; the enterprise user need only log onto their home system. Once logged on, the enterprise user can automatically access the users' files and services on Google without an additional login. Although convenient, this functionality increases the security exposure to not only the weakness of the enterprise system, but also to the weakness of Google's infrastructure. If, for example, Google's infrastructure has a security flaw, then it may be possible for someone in one enterprise to access accounts from another enterprise. On the other hand, security flaws in the enterprise system may lead to weaknesses in the access controls of the information managed by Google Apps. Additionally, connected applications may provide unintended connections among users, as was demonstrated with the introduction of Google Buzz. [17]

When each enterprise maintains its own infrastructure, a failure in one enterprise may cause failures across the cloud. Unless an enterprise uses a single cloud from a single vendor, integrating the various applications, infrastructures, and policies among many different clouds and cloud vendors will be a significant challenge. In fact, it will be a challenge to ensure that the different policies do not contradict and potentially permit access that should not be allowed at the system level.

Ultimately, the proof is in the pudding. Will the cloud vendors be willing to stand behind the security of their systems? In the case of Amazon's EC2 and Simple Storage Services (S3) services, Amazon suggests that their EC2 and S3 infrastructure not be used for systems that must satisfy the Payment Card Industry Security Standards [18], although it has published a paper on how Amazon Web Services can be used in a Health Insurance Portability and Accountability Act (HIPAA) compliant environment. [19]

In the HIPAA paper, Amazon essentially places almost all the requirements on the "user/enterprise"

to encrypt all the data stored and to manage its keys. Amazon provides services to log safely into its systems and provide some data recovery and integrity.

In the realm of reliability, prior to the breakup, AT&T was required to build systems that had an up-time reliability of "five nines" (about 5.2 min/yr downtime). Part of the reason for this was to ensure services in case of national emergency. Current cloud based systems are advertised as providing "three nines" (almost 9 hrs/yr downtime). [20]

Determining Where Trust Should be Placed

Clearly, there are many challenging security issues related to cloud computing. In our research, we are working on a formal, structured, possibly mathematical approach that will give users and cloud-developers deeper insight into what should be done, how it might be achieved, and where the trust should be placed. This research includes the investigation of implementation structures and assurance provisions for "security" in cloud-based systems. To do this, we will attempt to provide security architectures and models that satisfy the following:

- They are aware of the amorphous nature and scale of the cloud computing paradigm
- They include mathematical models of the security properties that can be used to help analyze those properties
- They provide the underpinnings on which applications/enterprise/user level security policies/properties can be implemented
- They provide the foundations on which the implementation assurances can be ascertained.

Our hope is that the results of the research will provide a framework that can be at least partially applied to the current cloud architectures and may lead to new architectures with better defined, more assured security.

Over the past 30-plus years in the operating system security world, a lot of work has been done to provide highly assured components with trustworthy systems. Unfortunately, the commercial world has ignored a lot of this work. Recent efforts have focused on the use of separation kernels. For example, Green Hills has recently received a National Information Assurance Partnership (NIAP) certificate for its Integrity 178B Separation Kernel. [21] Separation kernels provide a minimal set of operating system services on which other trusted services and applications could be built. These may be thought of as slightly more functional than a Virtual Machine Monitor (VMM), although Green Hills and others are looking to implement high assurance VMMs using their technology.

Our approach to the problem involves separation of 'virtual' resources. This approach constructs an infrastructure that establishes (or reconstructs where appropriate) resources, identifies and authenticates users, and then controls access to the resources. Our focus is to provide a model and a security architecture that provides the infrastructure that will accomplish these goals.

An Example

For instance, consider PaaS. An enterprise might wish to run its own applications. These applications may only run on an intermittent basis and/or require a large number of resources. One way to achieve this is to use a cloud PaaS.

We use the term 'enterprise' to describe the organization requiring the platform and 'provider' for the organization providing the cloud platform resources. The PaaS provider would provide 'platforms,' either 'real' as part of a virtual environment (a means for downloading an operating system and for managing the platforms), or as a possible network interface(s) on the platform(s). The enterprise loads operating systems, applications, *etc.,* onto the platform(s) and manages all the interfaces and resources provided. The example below assumes that multiple platforms will be used.

The security policy visible to the user includes:

- **Identification**—A set of platform names issued by the provider (unique to the enterprise)
- **Authentication**—A secure channel that can be used to load the operating system(s) onto the platforms—the provider is trusted to ensure that the only communication with the platforms is from or to the enterprise
- **Integrity**—The provider should guarantee that the resources are "empty" on first use and that none of the platform resources are modifiable by any party other than the enterprise. This includes any management functions; it is up to the enterprise to ensure that any network interfaces are appropriately protected
- **Privacy**—The provider should guarantee that there is no third party access to the platform processor, memory, and/or disk files
- **Provision of Service**—The provider should provide access to the resources on demand, per any service level agreements between the enterprise and the provider.

There at least two models of this kind of service:

1. Resources are provided on an ad hoc, intermittent basis. In this version, there is no connection between consecutive uses of the resources. The enterprise uses the resources once. During subsequent uses, the enterprise assumes that all the previous data does not exist or has been erased by the provider. The only connection between the two usages is that the enterprise uses the "same identifiers" to access new instances of the resources.

There is no guarantee that the same physical resources will be used for each run of the platform(s).

2. The enterprise 'turns off' the platform, but in subsequent use after turning it back on, finds the platform resources in the same state they were in after being turned off. As expected, the enterprise might pay more for this service. In this case, the provider must protect the information in the resources between runs from both modification and access by third parties. There is no guarantee that the same physical resources will be used in each run of the platform.

Note that in both cases, the provider provides access to platforms and associated data. The platforms are available to others when the enterprise is not using them. Any provider configuration data about the platforms must be protected from modification and, in the second case above, any enterprise information that will be reused must also be protected.

Informally, a portion of the model might then take the form of:

- **VPlatform**—The set of names of virtual platforms that will be provided to enterprises
- **VPlatformType**—Whether the VPlatform resources are persistent (type 2 above) or not
- **VPlatformResource**—The set of resources associated with a VPlatform
- **Enterprise**—The set of enterprises that use VPlatforms
- **Allocation**—An association of an Enterprise with a Platform, VPlatformType and VPlatformResources. The same Enterprise may have multiple VPlatforms, and VPlatformResources associated with it
- **PlatformCloud**—A sequence of sets of Allocations.

The security properties then become statements about the resources and platforms. For example:

No pair of allocations shares any common VPlatforms or VPlatformResources.

As depicted in Figure 1, the security properties can be modeled on a collection of the statements above. Each of the statements should map back to some aspect of the system's user-visible security property. We could use our statements about the relationships of the entities (sets) we describe to prove additional properties of the system.

Following the security model's construction, a high-level execution model should be constructed and validated mathematically to determine that it satisfies our security model. Next, it is necessary to map our high-level model to varied cloud aspect implementations as documented by the vendors.

Conclusion

Cloud security is an ill-defined, little-understood area of distributed computing. However, we believe that progress can be made to provide a level of assurance that accommodates the resources needed to support DoD and the federal government's information processing requirements. ■

About the Authors

Dr. Bret Michael | is a Professor of Computer Science and Electrical Engineering at the Naval Postgraduate School. He conducts research on the reliability, safety, and security of distributed systems. He is an Associate Editor-in-Chief of IEEE Security & Privacy magazine and a member of the IATAC Steering Committee.

Dr. George Dinolt | is a Professor of Practice in Cyber Operations at the Naval Postgraduate School. His research interests are primarily in the high assurance portions of Computer Security. His research covers formal methods and the connections between them and security policies, secure systems architectures and secure-systems design.

References

1. IAnewsletter, vol. 13, no. 1, winter 2010, p. 34.
2. Ibid.
3. M. Campbell-Kelly. "The Rise, Fall, and Resurrection of Software as a Service: A Look at the Volatile History of Remote Computing and Online Software," Communications of the ACM, vol. 52, no. 5, pp. 28–30, May 2009.
4. B. Michael. "In Clouds Shall We Trust," IEEE Security & Privacy, vol. 7, no. 5, p. 3, September/October 2009.
5. M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. "Above the Clouds: A Berkeley View of Cloud Computing," EECS Department University of California, Berkeley. Technical Report UCB/EECS-2009-28, 10 February 2009, *http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html.*
6. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Version 15, 7 October 2009, *http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.*
7. *http://en.wikipedia.org/wiki/Cloud_computing.*
8. *http://aws.amazon.com.*
9. *http://docs.google.com.*
10. *http://www.disa.mil/race*
11. H. G. Miller and J. Veiga. "Cloud Computing: Will Commodity Services Benefit Users Long Term? IEEE ITPro, vol. 11, no. 6, p. 67-69, November/December 2009.
12. *http://www.opencloudmanifesto.org.*
13. *http://www.fcw.com/Articles/2009/04/16/Cloud-computing-moving-into-public-safety-realm.aspx.*
14. *http://www.cloudsecurityalliance.org.*
15. *http://www.cloudsecurityalliance.org/csaguide.pdf.*
16. *http://www.google.com/apps.*
17. *http://www.nytimes.com/2010/02/15/technology/internet/15google.html.*
18. *http://www.mckeay.net/2009/08/14/cannot-achieve-pci-compliance-with-amazon-ec2s3*
19. *http://awsmedia.s3.amazonaws.com/AWS_HIPAA_Whitepaper_Final.pdf.*
20. *http://www.google.com/apps/intl/en/business/infrastructure_security.html.*
21. *http://www.niap-ccevs.org/cc-scheme/st/vid10119/maint200*

Pennsylvania State University

by Angela Orebaugh

In 1855, Pennsylvania State University (Penn State) was originally founded on 200 acres in Centre County, Pennsylvania, as an agricultural school that applied scientific principles to farming. Engineering Studies were introduced in 1882, making Penn State one of the nation's ten largest undergraduate engineering schools. Today, Penn State has grown into a large, geographically dispersed, major research institution. Nineteen campuses, 15 colleges, and one online World Campus currently comprise Penn State. In Fall 2009, Penn State served over 80,000 undergraduates and over 13,000 graduate students, with half of the student population enrolled at the main campus in University Park.

The National Security Agency (NSA) and the Department of Homeland Security (DHS) have designated Penn State as a National Center of Academic Excellence in Information Assurance Education (CAE/IA) since 2003 and National Center of Academic Excellence in Information Assurance Research (CAE-R) for 2008-2013.

The College of Information Sciences and Technology (IST) offers a bachelor's degree in Security and Risk Analysis (SRA). This degree program is intended to familiarize students with the general frameworks and multidisciplinary theories that define the area of security and related risk analyses. Courses in the major engage students in the challenges and problems associated with assuring information confidentiality, integrity (*e.g.*, social, economic, technology-related, and policy issues), as well as the strengths and weaknesses of various methods for assessing and mitigating associated risk. The major provides grounding in the analysis and modeling efforts used in information search, visualization, and creative problem solving. This knowledge is supplemented through an examination of the legal, ethical, and regulatory issues related to security that includes analyzing privacy laws, internal control, regulatory policies, as well as basic investigative processes and principles. Such understanding is applied to venues that include transnational terrorism, cyber crimes, financial fraud, risk mitigation, and security and crisis management. It also includes overviews of the information technology that plays a critical role in identifying, preventing, and responding to security-related events.

IST also offers a graduate degree in Security Informatics, which seeks to improve the cyber security of individuals and organizations by creating innovative solutions for detecting and removing cyber threats, recovering from cyber attacks, protecting privacy, enhancing trust, and mitigating risks.

Penn State includes a number of research centers focused in cyber and information security:

- The **Center for Information Assurance** plans, coordinates, and promotes IA research, education, and outreach. The faculty coordinators for the center include Dr. Chao-Hsien Chu and Dr. Peng Liu. The center's missions are:
  - Conduct broad-based research on various aspects (theoretical and applied; technical and managerial; wired and wireless, *etc.*) of information and cyber security
  - Educate and train information security professionals through degree and continuing education programs, and to insure that information security awareness is instilled in all Penn State students
  - Provide assistance and technical support to industry, non-profit organizations, government, and individuals in the information and cyber security area. [1]

- The **Networking and Security Research Center (NSRC)** was established in 2003 to provide a research and education community for professors, students, and industry collaborators interested in networking and security. It also provides a unique avenue for interaction with industry; the

▷▷ *continued on page 15*

Cloud Computing for the Federal Community

by Hannah Wald

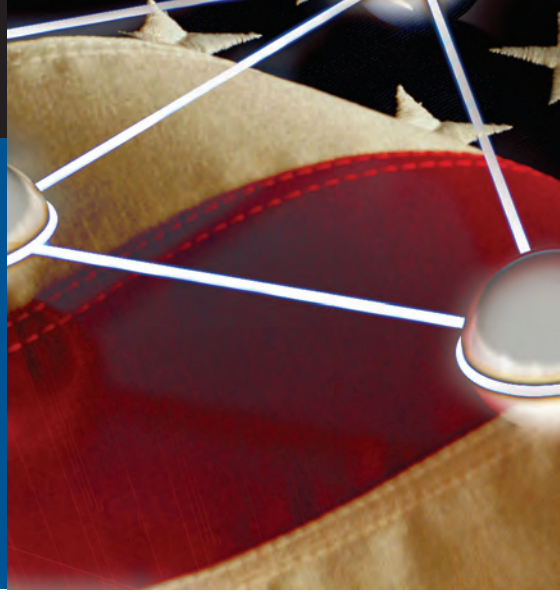

The question is not *whether,* but *when,* the U.S. federal government will embrace cloud computing. The current administration—in particular its Chief Information Officer, Vivek Kundra—is very enthusiastic about this technology's potential. Some federal agencies are already moving into the cloud: the Defense Information Systems Agency (DISA) is pilot testing a cloud [1]; the National Aeronautics and Space Administration (NASA) has announced plans to develop a cloud that can be used both internally and for collaboration with external research partners; [2] the Department of the Interior (DOI) has an Infrastructure as a Service (IaaS) offering called the National Business Center Grid (NBCGrid), with other offerings set to roll out in the near future; [3] and the General Services Administration (GSA) offers access to various externally provided cloud applications through its portal site, *http://apps.gov.* [4]

The federal government is not seriously considering cloud computing simply because of its hype. Agencies are finding it increasingly costly and difficult to procure, set up, maintain, and secure traditional computing architectures. This may explain why bodies such as the National Institute of Standards and Technology (NIST) and the Government Accountability Office are holding off on setting rules and standards for cloud computing while they survey the landscape and take an inventory of best practices. They are concerned about the risks inherent in cloud computing but do not want to restrict innovation. Pro-cloud civil servants believe cloud computing can make federal Information Technology (IT) and services cheaper, easier, and more secure—and it can—provided the cloud is implemented and managed properly.

> "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (*e.g.,* networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

For many federal agencies, a *community cloud* would be the best service model to use (regardless of the exact type of service being provided). The GSA, or another provider who is familiar with federal IT needs, could stand up a multi-agency cloud that facilitates and enforces compliance with government-wide security standards such as those outlined in regulations (*i.e.,* Federal Information Security Management Act [FISMA]) or guidance documents (*i.e.,* the NIST 800 series). Alternatively, individual cabinet-level agencies could provide clouds for their "community" of internal divisions, which could serve agencies' individual compliance needs more easily than a generalized multi-agency cloud. [5] DISA's Rapid Access Computing Environment sets a precedent for this model: it is intended to serve the entire Department of Defense, which has its own set of security standards in addition to those mandated for civilian agencies. [6] A third possibility is a "federated" hybrid of agency-specific community clouds and a government-wide community cloud, all with certain common standards (*i.e.,* minimal security baseline, universal protocols) but otherwise tailored to specific purposes.

Understanding the merits of a community cloud requires understanding fundamental cloud

computing concepts, starting with the definition of "cloud computing" provided by NIST:

"Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (*e.g.,* networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." [7]

NIST also lists five essential characteristics of cloud computing:

- **On-demand self-service**—A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.
- **Broad network access**—Capabilities are available over the network and accessed through standard mechanisms that promote use by client platforms (*e.g.,* mobile phones, laptops, and PDAs).
- **Resource pooling**—The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. A sense of location independence exists because the customer generally has no control over or knowledge of the provided resources' exact location but may be able to specify location at a higher level of abstraction (*e.g.,* country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.
- **Rapid elasticity**—Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear unlimited and can be purchased in any quantity at any time.
- **Measured service**—Cloud systems automatically control and optimize resource use by leveraging a metering capability appropriate to the type of service (*e.g.,* storage, processing, bandwidth, and active user accounts). The provider and consumer can monitor, control, and report resource usage, thus providing transparency of the utilized service. [8]

Industry expert Dave Linthicum, notes that cloud computing is similar to time-sharing on mainframes, but with some added features. For example, cloud clients can "mix and match" solutions using a software offering from one provider and an infrastructure offering from another. Commoditization of bandwidth allows clients to easily leverage distantly located resources—something that was difficult only a few years ago—and pay for use of those resources as if they were gas or electricity. Finally, cloud providers are particularly innovative in the services they offer and are developing new services all the time. [9] Cloud allows users to leverage IT solutions with an unprecedented level of granularity.

An organization can pay an outside cloud provider for data, applications, operating platforms, raw digital storage, and/or processing resources: Data as a Service (DaaS), Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS), respectively. [10] A data-mining company providing its customers with on-demand access to its records of individual purchase histories is an example of DaaS; Google Apps are SaaS; a firm offering application development environments to startups is selling PaaS; and a company offering access to raw computing resources is selling IaaS.

The split of assurance responsibilities between the provider and client varies depending on the service. With DaaS and SaaS, the provider has control over almost everything. With PaaS, the client is responsible for application security, and

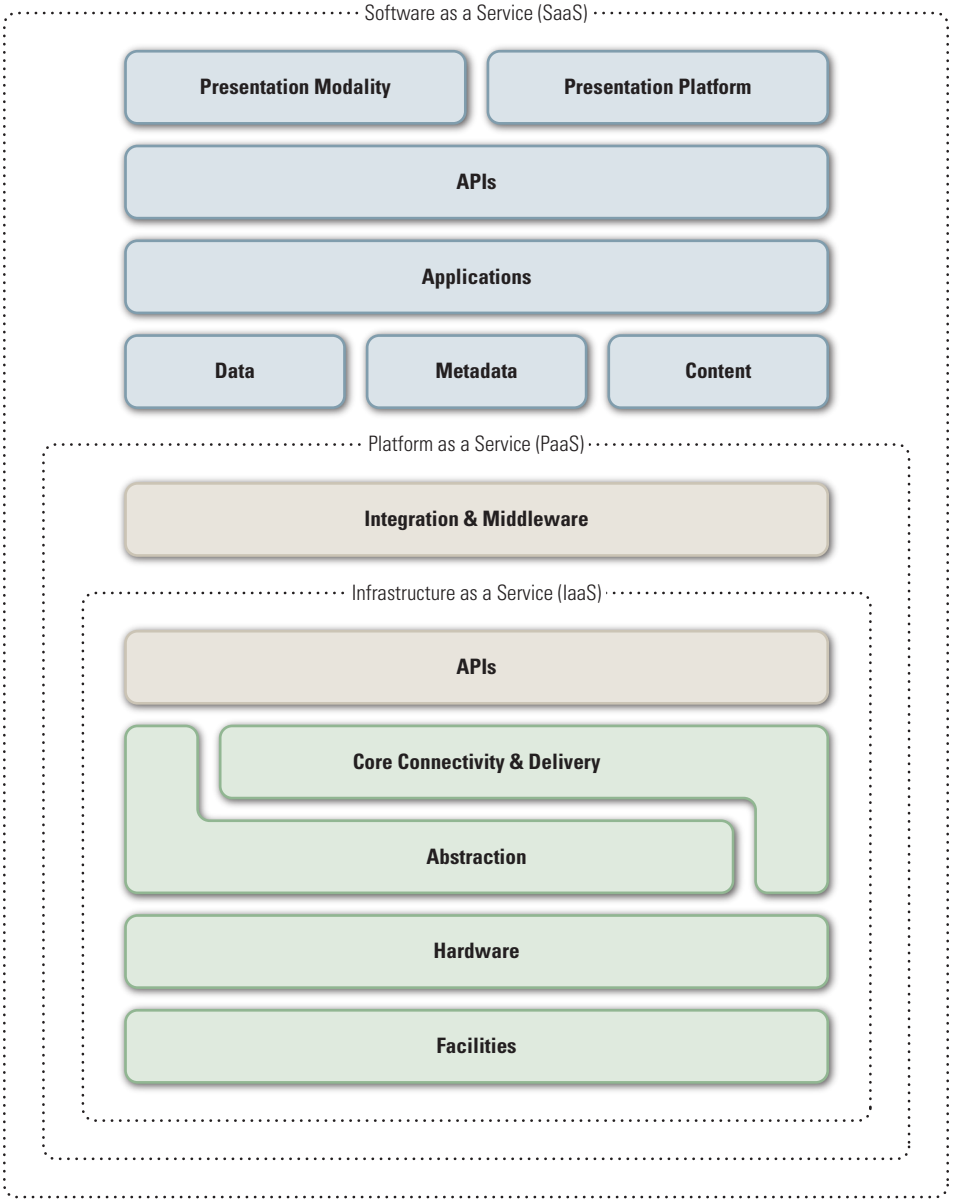

Figure 1 Provider Assurance Responsibility in Different Types of Service [11]

everything else is left to the provider. With IaaS, the client is responsible for everything but physical and (some aspects of) network security. Regardless of the service and inherent allocation of responsibility, cloud clients ultimately leave the fate of their information assets in the provider's hands (see Figure 1).

The service provider is responsible for maintaining, upgrading, and securing the hardware and software (where applicable) on which the service runs. Ideally, this setup allows users to stop worrying about the security of their information assets by leaving them in more competent hands. Cloud computing also has certain security advantages. For example, a desktop computer almost never complies with an organization's security policy "out of the box," but a cloud can be configured so every new virtual machine created therein is compliant. Monitoring certain activities and rolling out updates across a cloud is relatively easy—unlike doing so across a collection of distinct physical machines.

However, cloud computing presents a variety of information assurance (IA) challenges. One salient feature of the time-sharing model was trust. The users and owners of the old mainframes were part of a community with common incentives and goals, which is not necessarily the case in cloud computing. In a public cloud, the relationship between clients and providers is largely transactional, and the clients do not know each other. The parties involved have little basis for trust and may in fact *distrust* one another to a certain extent.

Trust, or lack thereof, is a factor in all five of the fundamental cloud security challenges. These challenges all involve uncertainties about the provider's standard of care and how the provider will treat the client (and the client's data) in the event of a problem. [12]

- **Data protection**
  - Where do data physically reside, and does the data's location have legal ramifications?
  - Are data safely protected (*i.e.,* by encryption) while stationary or in motion within and across the cloud?
  - How is availability of data assured in the cloud?
  - Does the provider take measures to ensure that deleted data is not recoverable?
- **Security control**
  - What security controls does the cloud provider need to implement, and how?
  - How are assurance levels effectively and efficiently managed in the cloud?
- **Compliance**
  - Is the cloud complying with all the necessary guidance?
  - Can the provider substantiate claims that security controls are implemented sufficiently?
- **Multi-tenancy**
  - Are my assets vulnerable if another client is exploited by an attack?
  - How does the cloud provider keep different clients' data separated and inaccessible from other clients?
  - If a forensic/electronic discovery procedure is conducted on one client's data, how will the

provider protect the confidentiality of other clients' data?

- **Security governance**
  - Who owns/accesses/deletes/replicates data in the cloud?
  - How can the client ensure policy enforcement?
  - How can the client measure and track service/network performance?

Figure 2 illustrates the layers of the cloud and associated layers of security.

Exacerbating these problems is the fact that contracts with public cloud providers almost always take the form of non-negotiable service-level agreements (SLA) that severely limit, at best, the client's ability to see, audit, or control back-end operations in the cloud. A client's ability to do so would create more difficulties than most providers are willing to deal with. The provider may not want to answer questions about its security practices. Cloud SLAs also generally absolve the provider of liability in the event of a security breach. (This is not the case with private and community clouds: more on this later.)

If the transition of federal information assets into the Cloud Computing Environment (CCE) is inevitable, then how can the federal government effectively mitigate the risks inherent in the cloud? First, government organizations must decide whether to move certain assets to the cloud at all. On the face of it, spending $10 a day for cloud infrastructure seems less costly than spending $100 on in-house infrastructure (not to mention capital expenditure; it is less costly to start up a virtual server in a cloud than to set up a physical one). However, thinking only in terms of $10 versus $100 for regular maintenance is dangerous because it ignores other kinds of costs. What will it cost an agency if moving to the cloud compromises its ability to protect sensitive data or meet mission requirements? Agencies need to consider these kinds of costs as they evaluate their information assets for "cloud readiness" on a case-by-case basis. [14] Once an agency has decided which assets it can safely transition to the cloud, it needs to choose the service model—or relationship between cloud client and provider—that best fits its requirements. The four cloud service models—public, private, community, and hybrid—have different sets of costs and benefits (see Figure 3).

The *public cloud* service model is probably what many people would consider the archetypal model of cloud computing. In the public cloud model, a provider sells cloud services to multiple unrelated clients, or tenants. They leave

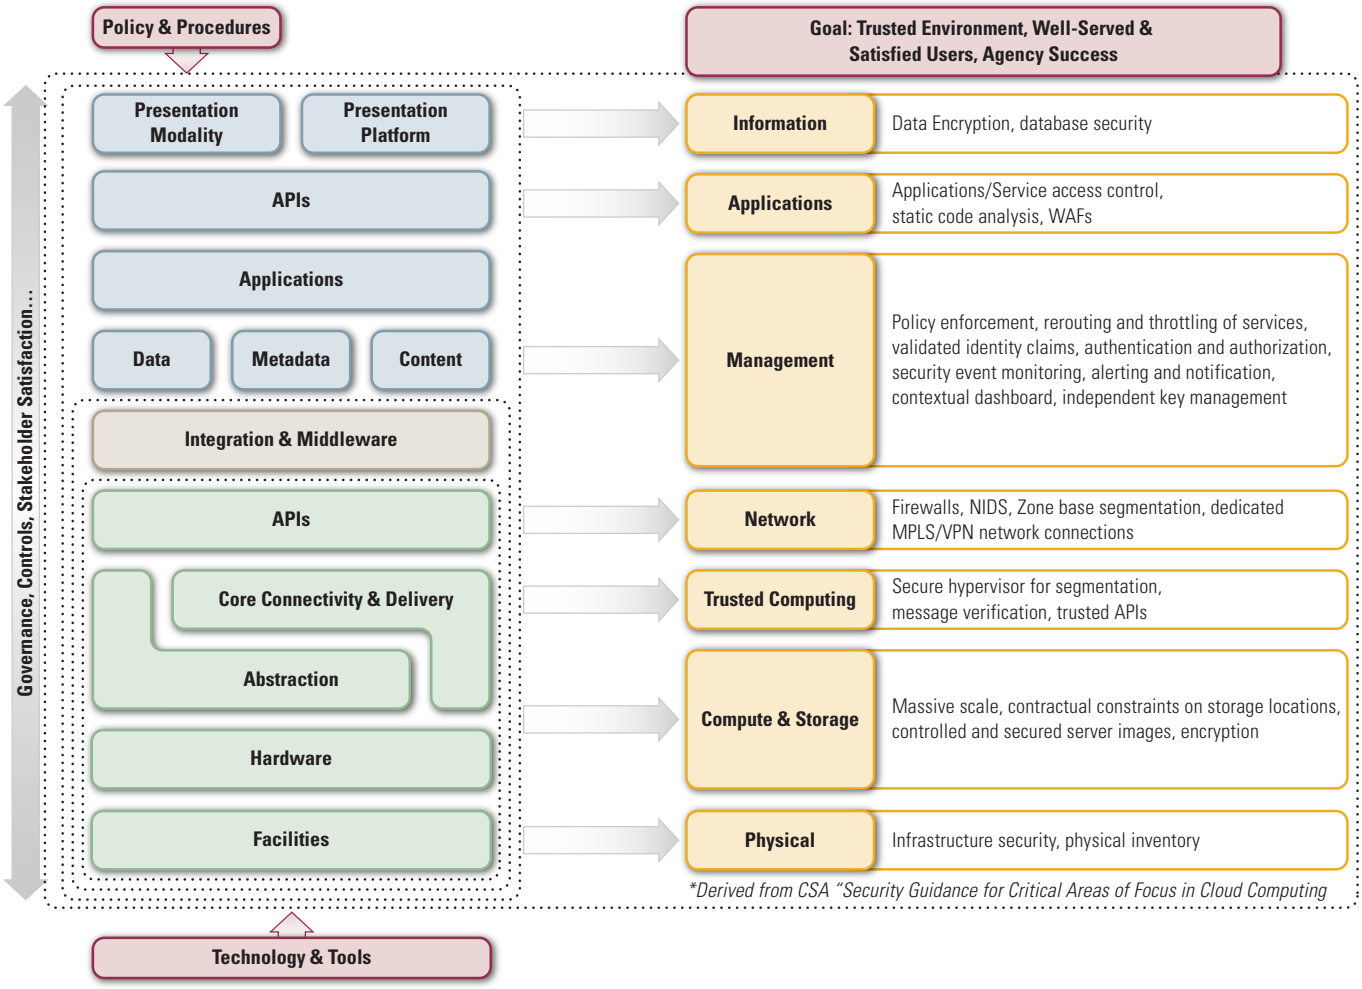

Figure 2 Layers of Cloud Computing Environment (CCE) Security [13]

back-end maintenance and operations to the cloud provider. This arrangement is very cost-effective and, in theory, lets clients rest easy knowing the security of their information assets is in good hands. However, the fundamental cloud security challenges mentioned earlier are most problematic in this model.

If a federal agency were to entrust its information assets to a cloud provider under the terms of a standard cloud SLA, the agency would have difficulty demonstrating compliance with IA standards mandated by regulations, such as the FISMA. Most public cloud providers would have to significantly retool their operations to help federal agencies meet their IA obligations. Some providers are attempting to do so (Amazon's "virtual private cloud" is an example [16]), but, for the time being, public clouds are inappropriate for anything but the least critical, most low-risk federal information assets.

A *private cloud* can be operated by the same organization that uses it, or a dedicated provider can operate the cloud on the organization's behalf. A private cloud, when managed properly, is the most secure type of cloud service model because it is directly controlled by its client. Private clouds also make more efficient use of physical IT assets than traditional data centers, but lack most of the economic benefits of outsourced cloud service. For organizations with less sensitive assets, putting everything in a private cloud may create unnecessary costs, inefficiencies, and redundancy. Also, if an organization has difficulty securing its information assets in a traditional setup, it is unlikely that transitioning to a private cloud will solve its problems. Such an organization would benefit from having a trusted service provider perform these functions.

A *community cloud* is somewhere on the continuum between the public and private service models, and it enjoys some of the benefits of both. Like a public cloud, community clouds serve multiple tenants. The difference is that the tenants are not strangers but related entities that share common characteristics and needs. An individual client community member, multiple members working cooperatively, or a dedicated provider can operate community clouds. Unlike public clouds, community clouds are built and operated on the clients' terms: they can be constructed to facilitate compliance with standards that all clients use. Of all the cloud models, the community cloud is most similar to time-sharing in terms of the level of trust between all stakeholders. This type of cloud also offers many of the economic advantages of the public cloud because it eliminates a considerable amount of redundant effort and cost. Members of the client community can pay the provider for only what they use, or for the utility and subscription cost. The latter would still likely total less than what the client would have paid to operate its own individual data centers.

The last type of service model is a *hybrid cloud*, which combines two or more of the service models described above. An organization could, for example, keep sensitive proprietary data in its own private cloud and collaborate on projects with industry partners in a community cloud. For users belonging to the organization, these two clouds would, in effect, be seamlessly integrated through a single sign-on system. The problem with hybrid clouds is that they share vulnerabilities in the system's least secure areas and present new vulnerabilities. For instance, if it is easy for a user to switch between clouds on his or her desktop computer, it is also easy for that user to make a mistake and expose sensitive data. In addition, integrated clouds mean integrated complex systems, which by definition are rife with potential security vulnerabilities.

Returning to the central point of this article, a federal community cloud can provide a guaranteed IA baseline for its clients, whether they are departments within an agency or multiple agencies. It can reduce the cost of providing effective security and eliminate significant redundancy. It can also be fully accountable to its clients and their oversight bodies (*i.e.,* Office of Management and Budget, Congress). The clients and their oversight bodies can have a reasonable level of visibility into, and control over, cloud operations. All primary stakeholders could work together to set policy and address problems. Last but not least, federal community clouds can be used to facilitate intra- and inter-agency cooperation within the framework of the Federal Enterprise Architecture.

Setting up a community cloud and governance structure that will

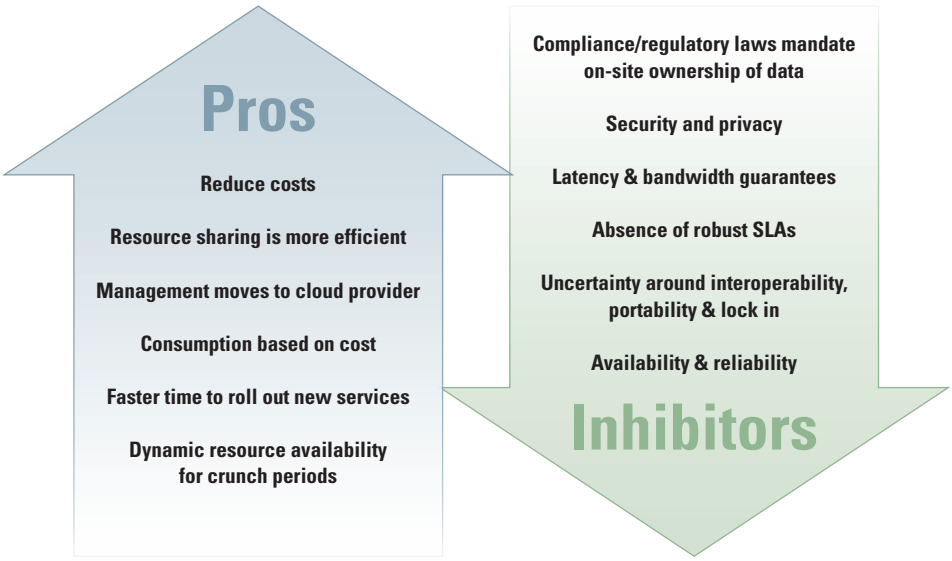

Figure 3 Advantages and Disadvantages of Cloud Computing From a Federal Perspective [15]

adequately satisfy all federal clients will be a challenging endeavor—even if the community is limited to the departments of a single agency. Architecting the technical and governance structure of a (possibly federated) community cloud for multiple agencies is an even more daunting prospect. A series of intra-agency (as opposed to inter-agency) community clouds may be the best possible outcome. Whether it serves only one agency or many, a community cloud is the most secure way for the federal government to realize the potential of cloud computing. ■

About the Author

Hannah Wald | is an Assurance and Resiliency consultant currently supporting the National Telecommunications and Information Administration at the Department of Commerce. Ms. Wald has contributed to the research conducted for IATAC's State of the Art Report on Supply Chain Security, which is scheduled for release this spring. This article draws heavily on research conducted and materials produced by her colleagues. Ms. Wald has a Master's degree in information science from the School of Information at the University of Michigan.

References

1. *http://www.disa.mil/race*
2. *http://nebula.nasa.gov*
3. *http://cloud.nbc.gov*
4. *https://apps.gov/cloud/advantage/main/start_page.do.* A link to a cloud service on apps.gov does not mean that the service is "safe" or that its provider has demonstrated compliance with federal security standards.
5. Some large agencies that are not at the Cabinet level, such as the Internal Revenue Service or Social Security Administration, may also benefit from having their own community cloud (admittedly, at that level the distinction between "community" and "private" cloud is not very clear).
6. On that note, some federal government entities—particularly those involved in law enforcement, defense, and intelligence—will need private clouds to protect their classified information assets.
7. Grance, Tim, and Peter Mell. "The NIST Definition of Cloud Computing." National Institute of Standards and Technology: Information Technology Laboratory Website. 7 October 2009. National Institute of Standards and Technology, Information Technology Laboratory, Web. Accessed 12 January 2010. *http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.* Page 1.
8. Ibid.
9. Linthicum, David S. Cloud Computing and SOA Convergence in Your Enterprise. Boston: Pearson Education, Inc., 2010. Pages 25–26. Print.
10. NIST's definition of cloud computing recognizes SaaS, PaaS and IaaS, but not DaaS. However, I have included DaaS because it is a fairly common cloud service offering.
11. Graphic from Hanna, Steve. "Cloud Computing: Finding the Silver Lining." 18 March 2009.
12. For a more in-depth discussion of security and legal issues in Cloud Computing, refer to guidance from the Cloud Security Alliance at *http://www.cloudsecurityalliance.org*
13. Graphic from Theodore Winograd, Holly Lynne Schmidt, Kristy Mosteller, and Karen Goertzel, "Public Cloud Computing Environment (CCE) Acquisition: Managing Risks to the Federal Government." Booz Allen Hamilton, 2009.
14. Linthicum 2010, pp. 192–193.
15. Graphic from Stephen T. Whitlock, "Cloud's Illusions: Jericho Forum Future Direction." 16 February 2009.
16. *http://aws.amazon.com/vpc*

▷ *continued from page 9*

IATAC SPOTLIGHT ON A UNIVERSITY

members of the NSRC actively consult with industry and participate as partners on funded projects. Member companies enjoy benefits for sponsoring research and having access to the latest results and technical reports from the NSRC. Hosted in the Department of Computer Science and Engineering (CSE) at Penn State, the NSRC is comprised of nine faculty members in the College of Engineering, including eight members from CSE and one from Electrical Engineering (EE). Several faculty members also have joint appointments in EE and the College of Information Sciences and Technology. The NSRC includes approximately 50 Doctor of Philosophy (Ph.D.) and Master of Science (M.S.) students, and several undergraduate honors theses are advised through NSRC faculty as well. [2]

▶ The **LIONS Center** is the IST Center for Cyber-Security, Information Privacy, and Trust whose mission is to:

- Detect and remove threats of information misuse to the human society: mitigate risk, reduce uncertainty, and enhance predictability and trust
- Produce leading scholars in interdisciplinary cyber-security research
- Become a national leader in information assurance education.

The center currently includes seven core faculty members, 20 collaborating faculty, two research associates, and 19 Ph.D. students. The center has published over 200 publications since 2002 and received over $3 million in research grants. ■

References

1. *http://net1.ist.psu.edu/cica/cia-ist.htm*
2. *http://nsrc.cse.psu.edu*

DISA RACE: Certification and Accreditation for the Cloud

by Munjeet Singh and Troy Giefer

Background

Since the Obama Administration announced plans to use cloud computing to cut costs on infrastructure and improve performance of government computing systems, the Department of Defense (DoD) and other federal agencies have become increasingly interested in how to take full advantage of the potential benefits offered by cloud computing. [1] Few existing cloud providers meet DoD requirements and choices are primarily limited to the public domain. Additionally, there are concerns about government use of public clouds because of the lack of control and visibility into the cloud's underlying security infrastructure and the challenges of complying with DoD and federal information assurance (IA) policy and procedures.

Given the high level of interest in cloud computing, the Defense Information Systems Agency (DISA) recognized the need for a government-managed cloud that could benefit the DoD community. DISA subsequently developed the Rapid Access Computing Environment (RACE), which is an agile and robust cloud computing environment that allows DoD organizations to provision virtual servers and storage from a Web portal. RACE is a streamlined workflow process designed for use in a virtualized development and test environment. RACE is customized to enable DoD components to rapidly and seamlessly transition from application development to testing and into a full production environment, a process known as the Path-to-Production. Current DoD certification and accreditation (C&A) policy has been adapted to allow organizations to use RACE cloud resources, thereby quickly connecting to the cloud while complying with DoD requirements.

This article describes the goals DISA sought to achieve and the approach it took as it developed the RACE Path-to-Production process. It will also highlight many of the key characteristics and capabilities of the DISA RACE cloud.

Goals and Objectives

DISA's primary goals in developing the RACE Path-to-Production were to:

- Develop a streamlined C&A process that would reduce time and effort required to transition an application from development to test, and ultimately to a production environment (Path-to-Production process)
- Reduce the current C&A approval time from 120 days to under 40 days
- Develop an enhanced RACE portal that enables customers to purchase and manage virtualized RACE development and test environments and provided additional storage.

Approach

Before designing a new streamlined C&A workflow process, it was important to understand the current approval process, identify key organizations involved in the decision making, and identify the artifacts required by each organization. The approach used in developing the Path-to-Production process was conducted in two phases.

Phase I consisted of data gathering and documentation of the current C&A workflow process. This included identifying all key organizations involved in data collection, document handling and processing, validation, certification, and accreditation of a system. Personnel from each organization involved in the approval process were interviewed to define roles and responsibilities. The responsibilities of each entity were then mapped to a process flow diagram that identified each step in the process. In addition, a complete list of artifacts required by each key organization as input and generated as output was compiled. The end result captured the comprehensive 'as-is' DoD Information Assurance Certification and Accreditation Process (DIACAP). DISA supplemented process steps required to obtain certification.

Phase II consisted of a duplication analysis of the organizational roles and artifacts. The intent of the analysis was two-fold, specifically to: (1) eliminate duplication of effort across the various organizations involved in the C&A

workflow process; and (2) reduce or eliminate duplication of documentation. Eliminating duplication of effort across the organizations involved in the decision making would reduce the time required for a system to reach approval to operate (ATO). In addition, eliminating the duplicate documentation would both reduce the possibility of inconsistencies and eliminate the need for the customer to create multiple documents that contain duplicate information, which would further reduce the time to complete the C&A process.

The analysis of the current processes, responsibilities, and artifacts gave DISA the groundwork for designing a more efficient C&A workflow process (Path-to-Production).

Path-to-Production

DoD organizations use the RACE cloud for application development and testing, and to prepare for deployment into a production environment. Path-to-Production refers to the process that an organization follows to promote the application developed in a virtualized environment from development to test, and from test into a Defense Enterprise Computing Center (DECC) production environment (Figure 1). The Path-to-Production process reduces the total time required to obtain accreditation of an application from an average of 120 days to under 40 days, in part, by streamlining approval workflows and leveraging inheritance of IA controls from the RACE cloud and DECC environments.

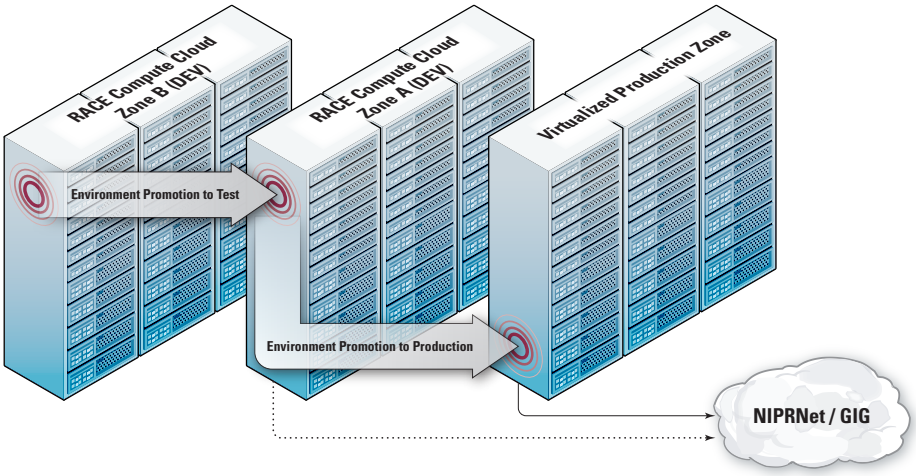

Figure 1 Path-To-Production

A number of characteristics were incorporated into the RACE Path-to-Production process that were key to streamlining and customizing the current process. DISA focused on the areas that offered the greatest return:

- Define standards and entrance criteria
- Streamline the approval process
- Reduce or eliminate duplication of effort and documentation
- Incorporate inheritance of IA controls as defined by DoDI 8510.01
- Develop hardened virtual operating environments (VOE)
- Implement a RACE portal.

RACE Standards

A key aspect of designing the RACE Path-to-Production process was defining a set of standards that provide the framework of the streamlined process. These standards enable rapid provisioning and promotion within the virtual environments. Examples of RACE standards include:

- The development and test process must be completed in a virtualized environment.
- Customers must start with provisioned VOEs provided by RACE.
- The Enterprise Mission Assurance Support Service (eMASS) application must be used as the C&A automation tool and central repository.

- Customers must adhere to the RACE standard set of ports and protocols while in development, test, and production environments.
- Vulnerability Management System (VMS) must be used to track asset-level vulnerabilities.
- A minimum of an Interim Approval to Test (IATT) is required to move an application into the RACE Testing environment.
- An IATT is valid for 90 days while in the test environment.
- A minimum of an Interim Approval to Operate (IATO) is required to move an application into the DECC production environment.

Recognizing that organizations often have unique needs that may fall outside of the standards established by RACE, DISA developed an exception resolution process to facilitate discussions between a RACE representative and the RACE customer to determine a resolution.

Streamlined Approval Process

Delegation of approval responsibilities to the lowest organizational level possible was key to streamlining the RACE C&A approval process. This approach resulted in a more agile workflow adaptable to the robust environment of the RACE cloud. To facilitate this streamlined approval process, the DISA Chief Information Officer implemented an Information Assurance Manager (IAM) role created specifically to manage activities within the RACE cloud. The RACE IAM's primary role is to provide a final review and approval of the application and virtual environment before it is promoted to the test and production environments. The IAM reviews the RACE customer's documentation to validate the accreditation decision made by the customer's Designated Approval Authority (DAA). In addition, the IAM considers additional application-specific data such as the ports, protocols, and services used by the system or application, and the proposed network topology. The RACE IAM also conducts joint validation activities of the IA controls with the customer early in the process, and establishes the parent/child inheritance relationship, which allows the system to inherit IA controls from the RACE cloud. This early coordination activity between RACE customers and the RACE IAM supports users as they move through the Path-to-Production process, ensuring that potential challenges are addressed early in the process.

The RACE C&A approval process is a joint effort shared between the RACE IAM and the customer. The customer conducting application development in the RACE cloud has the primary responsibility to oversee the validation, certification, and accreditation of the system or application as it progresses through the Path-to-Production process.

Duplication Analysis

The duplication analysis of the existing C&A approach and workflow process revealed more opportunities to streamline this process. The team identified opportunities to reduce the amount of documentation required for a successful accreditation. At each approval level, organizations had developed unique checklists of requirements and artifacts. This often required the customer to duplicate data in multiple documents. Further analysis revealed that a number of documents could be eliminated because the data was available in other C&A artifacts. Elimination of such duplication significantly reduced the time and effort spent on developing and reviewing C&A artifacts.

DISA implemented a key tool—eMASS—within RACE to manage the C&A workflow and documentation. A government-owned solution, eMASS integrates several capability models to support IA program management needs. It allows an organization to enter system information and to track the progress of information assurance activities (such as validation procedures, compliance status, and attachments) and associated action plans for sharing system security information and compliance status.

Inheritance of IA controls

Inheritance of IA controls was also key to streamlining the Path-to-Production process. RACE customers can directly inherit IA controls from the RACE cloud and DISA DECC (Figure 2). DoDI 8500.2 defines 32 controls that an automated information system (AIS) may inherit

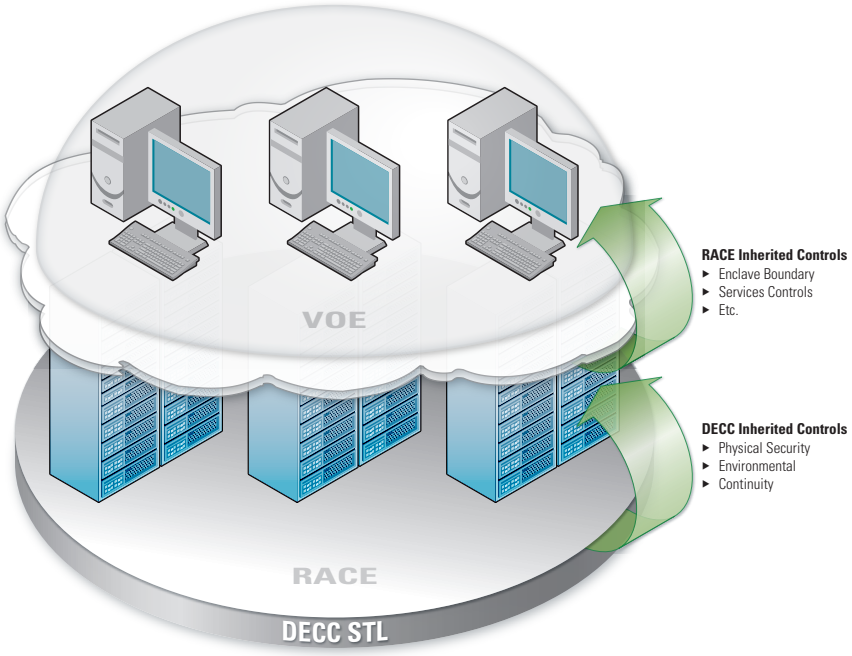

Figure 2 IA Control Inheritance

from the enclave in which it resides. The implementation, validation, and monitoring of these controls are the responsibility of the enclave and not the AIS. RACE customers inherit these controls, as well as the status and artifacts associated with the validation of each control.

This automated inheritance of IA controls is defined within the eMASS application. RACE serves as the parent system for a parent-child inheritance relationship used for all registered systems within eMASS. Every application that a RACE customer registers within eMASS will automatically be set as a child to the parent (*i.e.,* RACE) enclave, establishing inheritance. A pre-determined list of DoDI 8500.2 IA controls is automatically set as inherited from the parent in every child. For example, physical security is the responsibility of the parent enclave, not the responsibility of the child.

Hardened Virtual Operating Environments

Virtual operating environments are provisioned to RACE customers for use in the development and test environments. The VOEs are delivered with a development-friendly Security Technical Implementation Guides (STIG) implementation, streamlining both the development process and the C&A process for RACE customers. RACE offers the available virtual operating environments, as listed in Table 1, which are in compliance with DoDI 8500.2 at the Mission Assurance Category (MAC) II-Sensitive level.

Operating System	Architecture
Windows Server 2003	32-bit
Windows Server 2003	64-bit
Red Hat 4.6	32-bit
Red Hat 4.6	64-bit
Red Hat 5.1	32-bit
Red Hat 5.1	64-bit

Table 1 RACE Virtual Operating Environments

DISA has configured the virtual images to be compliant with a variety of DISA STIGs, to include Windows Server 2003 operating system, UNIX, Internet Information Services (IIS), and database checklists. The DISA team reviewed the recommended security settings from these STIGs to determine which had the potential to restrict application development. The VOEs are provisioned to RACE customers with those particular security settings left in a 'non-compliant' status. This practice allows customers to begin development immediately and provisions a consistent development environment for all customers.

However, these security settings will remain in a 'non-compliant' status only in the RACE development environment. The RACE customer is responsible for properly configuring these security settings to achieve a compliant status before promoting the application to the testing and production environments.

The VOEs are also provisioned with the latest Information Assurance Vulnerability Management (IAVM) patches installed. Once the VOEs have been provisioned, the customer assumes responsibility for keeping the images patched.

RACE Portal

A key component of cloud computing is the ability to provision and maintain environments in a self-service portal. DISA Circuit Switched Data (CSD) has implemented this ability through an enhanced RACE portal that allows RACE customers to take control of their environments with respect to the following functions:

- Ordering development, test, and production virtual environments
- Ordering additional storage for an existing virtual environment
- Promoting the environment from development to test or test to production
- Archiving the environment to tape backup
- Restoring the environment from an archive.

In addition, the RACE portal provides a document library that includes all IA documentation that will be used throughout the Path-to-Production process.

On the Horizon

DISA CSD is continually seeking opportunities to improve the Path-to-Production process to make it even more agile. This includes implementing automation to further reduce the C&A burden on RACE customers, and strengthening the IA posture of VOEs *via* integration of Host Based Security System (HBSS) into the RACE enclave. For more information, visit *http://www.disa.mil/RACE* for the latest news. ■

About the Authors

Munjeet Singh | is an information assurance contractor consulting as the Project Manager and Lead Engineer on cloud focused initiatives in the DoD domain. He is currently involved in deploying cloud and data center optimization initiatives to clients in DISA and across the Army.

Troy Giefer, CISSP, | is an information assurance contractor consulting on cloud computing research and the development of cloud computing security solutions for the DoD marketplace. Troy is a key lead in the effort to customize DIACAP for use in the DISA RACE cloud.

References

1. *http://www.whitehouse.gov/omb/budget/fy2010/assets/crosscutting.pdf.*

Look Before You Leap:

Security Considerations in a Web 2.0 World

by Sara Estes Cohen and Shala Ann Byers

Introduction

In recent years, social media, also known as Web 2.0, has emerged as a popular and powerful technology that enables individuals to collaborate, communicate, and share information from anywhere and at anytime. Currently, more than 30% of the world's population visits Facebook.com on a daily basis [1], and approximately 22% use YouTube to watch online videos. [2] First established within the commercial industry, this technology made popular the economically savvy use of low-cost social media technology. The federal government has since followed suit, launching organizations and government agencies into the foray of social media as a way of connecting with the public.

On January 21, 2009, President Obama signed the *Memorandum on Transparency and Open Government,* encouraging agencies to "establish a system of transparency, public participation, and collaboration." [3] On December 8, 2009, the Director of the Office of Management and Budget (OMB) issued the *Open Government Directive,* providing guidelines and deadlines for all federal agencies on developing their own 'open government' programs fostering the principles of transparency, participation, and collaboration. [4]

Agencies like the Department of Justice, the Library of Congress, and the Department of State responded by establishing Facebook profiles to communicate with the public. Additionally, the Federal Bureau of Investigation started a Twitter account to send daily news updates to the public. The Centers for Disease Control and Prevention (CDC) posts weekly *Hurricane Health and Safety Tips* on its Web site and distributes them to registered users *via* e-mail, mobile phone text messages, and Twitter. [5]

While embracing social media is key to succeeding in a new communications environment, effective strategy, planning, and support before launching a social media program are equally important. The results of an unstructured and disorganized adoption of social media can have serious complications, including data leaks or breaches in security from which it can be difficult—if not impossible—to recover.

To avoid these complications, it is imperative for an organization to identify a 'best-fit' solution based on internal goals, requirements, and challenges, before launching a social media program. Most importantly, organizations must standardize how they implement social media and develop training to educate users. Finally, organizations must institute a mechanism to enforce security compliance to ensure the protection of the information shared within the social media platform.

Framework

There are generally three approaches for implementing social media:

- Internal
- External
- Hybrid.

Each approach differs in location and ownership of underlying infrastructure (*e.g.,* government or privately-owned), audience (employees, the public, or both), and direction of communication (within, outside of, or across the firewall):

- **Internal**—Technology and infrastructure sit behind a firewall and are owned by the organization. This model consists only of internal communications, information and data exchange, storage, and management (within the organization, not across the firewall) and requires development of organization-specific solutions, tools, and technology.
- **External**—This approach leverages public social media for specified applications. For example, existing social media sites (*e.g.,* Facebook and Twitter) may be used for constituent relations and outreach. This model requires extensive strategic planning to target the

appropriate user groups with the right information. Additionally, this model must include organization-wide standardization to ensure consistency with respect to messaging (content/brand), security practices, and access to public sites and tools from behind the firewall.

- **Hybrid**—This model uses internal solutions (behind the firewall), developed by the organization for internal communication and operations, while simultaneously leveraging external, public social media for outreach and general communications. Like the external model, the hybrid also requires standardization to ensure the security of personnel, data, and information.

This article focuses on security considerations and challenges associated with the hybrid model, as it is the most complex of the three types of approaches. Because of its reliance on both internal and external infrastructure, the hybrid model must adhere to both internal and external, organization-specific security, management, legal, and communications policies.

Strategic Planning

To begin, an organization must first identify its goals and objectives for adopting social media. Identifying appropriate budget, development time, specific features and functionalities required, and level of intended risk are all factors to consider before implementing a social media strategy; by doing so, organizations can avoid developing an ill-fitting program. The following section outlines and discusses several planning considerations to assist in establishing a 'best-fit' approach.

Audience

Who is your target audience? This question can be answered by first defining the organization's responsibilities. Are you required to communicate with your constituents? Will you need to communicate with your employees during a crisis, or on a daily basis? These answers will help the organization clearly define its purpose for using social media; identify the tools that can accomplish that purpose; and successfully engage its audience using social media. Identifying your audience can also help determine the most appropriate Web 2.0 model and the best tools and technology to use.

Technology and Applications

Organizations can leverage social media for many purposes, including daily operations, outreach and awareness, constituent communications, emergency management, and business continuity, among others. Additional applications may include training, alert and notification, employee accountability, situational awareness, information gathering, and emergency communications. As technology advances and user awareness improves, the potential for using social media will grow accordingly.

Social media is not just about the technology or the tools—it is also about what the technology can help users do. Organizations must leverage social media in a way that resonates best with the targeted community, chosen goals, and objectives.

Additionally, proactively identifying potential applications before choosing and implementing social media tools can help avoid the 'Shiny New Toy' syndrome—investing in a tool that nobody uses because it does not meet organizational needs. A strategic approach will help ensure that the program is functional—for both the audience and organization—while remaining aligned with the desired goals and objectives.

Standardization

Standardization is the most important aspect in adopting social media. Social media standards must be developed in line with both organization-specific and external information technology (IT),

security, communications, operations (management), and contractual/legal policies and requirements. Organizations must establish standards for how they implement their own social media solution; there is no one-size-fits all solution.

Without some form of centralized guidance, departments might develop policies and processes that are inconsistent across the organization as the popularity and use of social media grows. This situation could result in varying levels of security and inconsistent security procedures. To avoid this, the organization must establish technical requirements and training standards regarding how all departments and components may use internet-based capabilities. Additionally, the organization must establish and disseminate organization-specific policies and procedures regarding technical, legal/contractual, communications, and management concerns. Each department may have additional requirements but, at a minimum, its practices should comply with the organization-wide requirements.

Security Requirements

Security requirements must take into account several factors, such as:

- The purpose the social media is intended to accomplish
- How social media will be used (application)
- What type of information will be exchanged (*e.g.,* classified information, Sensitive But Unclassified [SBU] information, Personally Identifiable Information [PII]) and the associated handling requirements
- How and where data will be stored
- Criteria for accessing the information
- How exceptions are managed
- What technical support will be required
- How the factors above will be affected by organization-wide use of social media.

Each of these factors must be taken into consideration to develop suitable and sustainable standards essential for enforcing compliance.

Social Media Guidelines and Governance

Federal policies and guidance governing the use of new and emerging communications technologies, as well as industry best practices for social media, should be carefully evaluated and followed to ensure compliance. If an organization is just beginning its foray into social media, it should consider using *Guidelines for Secure Use of Social Media by Federal Departments and Agencies,* released by the Federal Chief Information Officers Council in September 2009, as a starting point. [6]

Agencies need not start from scratch however – the General Services Administration (GSA) has already contacted third-party providers Flickr, YouTube, Vimeo, and blip.tv to develop government-specific terms of service. Additionally, GSA determined that Twitter's standard terms of service are consistent with government use and thus need no additional changes. [7]

Additionally, organizations should consider drafting their own social media engagement guidelines before allowing unfettered access to social media and online communities. A great example is the Air Force's Web Posting Response Assessment Flow Chart V.2., which explains the Air Force's internal policy on blogs and how to handle both positive and negative commentary posted online. [8] Such guidelines not only protect the organization from a legal standpoint; they can also help employees understand the implications of personal use, and how to develop and maintain social media tools in a way that complies with the organization's standards and best practices.

Risk Management

It is no longer feasible to dismiss the use of social media entirely because of its potential risk. Web 2.0 users are tech-savvy and will continue to find new ways to access and use social media despite an organization's best efforts to ban the technology. Instead of banning social media outright, organizations should identify how to use social media safely and securely. As with adopting any new technology, risk assessment is an integral aspect of adopting social media and must be conducted on a regular basis, allowing for adjustments over time to accommodate changes in technology and the threat environment.

The decision to adopt social media should be based on a strong business case that considers an organization's mission, technical capabilities, threats, and the expected benefits of adopting this technology. For example, national security agencies must protect classified data, whereas agencies or organizations that handle PII must protect the privacy of individuals. Consequently, different organizations have different priorities for security and privacy, and must address those priorities accordingly.

Challenges

After identifying a 'best-fit' solution and socializing the standards, the organization must develop an implementation plan and provide the continuous, reliable support needed for maintaining a structurally sound and sustainable program. Throughout the development and implementation of a social media program—whether internal, external, or hybrid—organizations should consider and address the following challenges related to security, technology, and infrastructure.

Information Assurance and Operational Security

A social media strategy must incorporate information assurance and operational security (IA and OPSEC) policies and procedures—as well as an

organization-wide training, education, and awareness package—focusing on IA and OPSEC issues to ensure that the policies and procedures are followed. Otherwise, data leaks and OPSEC violations are more likely to promulgate across all forms of electronic communications, including e-mail, social media, and Web sites. The organization must also address policies and develop compliance measures regarding access control, authentication procedures, account and user management, encryption, content assurance, and general communications security (COMSEC).

The requirement to address IA and OPSEC is nothing new. Concerns about social media are essentially the same as those that arose with the proliferation of the Internet and e-mail. Communications policies and information security procedures that apply to social media are similar to those that have traditionally applied to other forms of communications—whether electronic communications (*e.g.,* e-mail) or more traditional forms of communications (*e.g.,* letter writing or meetings).

Privacy and Confidentiality

Federal departments and agencies are bound by privacy requirements based on the Fair Information Practice Principles (FIPP), which require rigorous controls and procedures to protect the privacy of individuals. PII includes any information that can be directly associated with an individual. Those organizations that collect PII must put policies and procedures in place to handle, store, and dispose of PII securely. Such measures may address terms of use, legal ownership of PII, and the consequences of using or disseminating PII inappropriately.

In addition to addressing privacy policies, organizations must also be aware of threats to privacy and must implement measures to ensure that privacy is maintained. For example, some social media protocols (*e.g.,* certain programming languages, social media etiquette, *etc.*) may place PII at risk of exposure. Once exposed, PII could place individuals at risk of identity theft and fraud. An organization can reduce this risk by implementing enhanced protection measures for sharing data in interconnected systems, implementing monitoring capabilities and protocols, and educating users on proper social media etiquette ("safe-surfing").

Despite these challenges, agencies and organizations dealing primarily with private, sensitive, or classified information are not necessarily precluded from adopting social media. Rejection of social media also poses risks; organizations that choose not to leverage social media and new technologies may become obsolete over time.

Furthermore, unless an organization bans access to social media completely (which is nearly impossible to do), employees will inevitably use social media from within the organization's network. Those organizations that do not establish policies regarding the use of social media, and do not implement processes to protect their infrastructures from unauthorized use of social media, expose themselves to serious legal and security-related problems. Both their information infrastructures and their reputations can be irreparably damaged.

Technical Support

Although social media may seem to offer a quick and efficient communications solution, it comes with some technical challenges:

- **Bandwidth**—Social media sites may require more bandwidth than traditional sites. Therefore, organizations may require additional network infrastructure to support wide-scale use of external, resource-intensive Web sites (*e.g.,* YouTube, Facebook, *etc.*). If the organization is successful in engaging its audience in using social media, user demand will increase dramatically, ultimately increasing demands on network infrastructure. Consequently, the social media functions may compete with the organization's other functions for use of the network, which could impair overall mission capabilities over time. Organizations must plan for and ensure adequate bandwidth is available for widespread Internet use. Most hosting environments can provide additional bandwidth to cover surges in Internet or network activity. Organizations should develop memorandums of understanding (MOU) with their respective hosting companies to ensure sufficient bandwidth is available during surges of activity that may occur due to emergency events, times of heightened network activity, and increasing popularity in social media.
- **Malicious Attacks**—To one extent or another, all networks are subject to malicious attacks. Use of social media may increase that risk because, as more external Web sites are accessed, malicious actors have more opportunities to access an organization's networks and operational data. Implementing security controls across all Web 2.0 servers and verifying that sufficiently rigorous security controls are in place can reduce the threats to internal networks and operational data. Additionally, separating Web 2.0 servers from other internal servers may further mitigate the threat of unauthorized access to information through social media tools and Web sites.
- **Network Monitoring**—Foreign intelligence services (FIS) have extensive resources and have repeatedly demonstrated their capability to use automated 'social engineering' techniques to mine social media sites. By their very nature, social media sites have an abundance of information, which makes them susceptible to data

mining. Our adversaries can use this data to analyze aggregated information. Without adequate network monitoring (and user education), an organization cannot ensure that users are complying with its policies regarding the release of high-value information. Additionally, programming languages used in Web 2.0 applications (*e.g.,* Java, Ajax, and the JSON data interchange format) may create other opportunities for malicious actors to access an organization's back-end network infrastructure and do irreparable damage (*e.g.,* access or corrupt data or applications). Consequently, an organization using social media may need to implement increased security controls for any separate sensitive information residing on the server's backend.

Compliance and Enforcement

User education and training have always been crucial in safeguarding networks and data. However, with the advent of social media, training programs must be augmented to address the additional risks posed by social media. As organizations develop and adopt social media, users must understand the severity and nature of potential threats to security associated with its use. Organizations can incorporate social media training into their annual security training programs and address social media tools and sites during existing certification and accreditation procedures, thereby helping to ensure that their security standards are upheld. Additionally, organizations can develop a social media mentoring program, leveraging the skills of those employees with more advanced social media skills to train those for whom this technology is unfamiliar.

Incident Response

Finally, despite best efforts to train users on 'safe-surfing' and develop safeguards for protecting data and information, incidents will inevitably occur. Organizations must plan and develop measures for quickly responding to and recovering from data spills, misinformation and rumors, and malicious attacks. An important aspect of handling social media is anticipating such incidents, then developing and implementing a plan for managing and responding to them. Such planning will help ensure that social media becomes an integral part in an organization's communications toolbox.

Conclusion

Trends in communications and technology are increasingly dynamic and fast-paced. To keep up, organizations in both the public and private sectors must readily adapt by developing social media capabilities of their own. Although embracing social media is imperative to succeeding in a new communications environment, doing so without adequate planning can do more harm than good.

Social media is not a one-size-fits-all solution. Each Web 2.0 tool has its own purpose, audience, and challenges that must be considered carefully. As with any tool, a Web 2.0 tool must be chosen, not based on popularity, but on how effectively it meets the organization's needs and selection criteria.

Finally, an organization's social media program must align with its goals, objectives, budget, desired features and applications, internal and external security, IT, legal, and communications policies and requirements. Once implemented, the program must be standardized across the organization through socialization, education, and consistent training. Compliance with these standards must be upheld through consistent enforcement; proactive engagement is crucial to the security of an organization's networks, infrastructure, information, audience, and reputation. With well-thought-out strategy, planning, policies, procedures, and technical support, organizations may successfully and securely leverage social media.

Thank you to DeZario Morales, Akira Ikuma, Matthew Doan, and Mark Macala for their contributions to this article. ■

About the Authors

Sara Estes Cohen | has ten years of experience in communications and three years specifically focused in emergency response, continuity of operations, business continuity, and critical infrastructure protection. For her masters thesis, "Using Social Networking for University Emergency Communications," Ms. Cohen worked with the University of California, Los Angeles (UCLA) to develop a model for universities to engage in social media for emergency communications. Ms. Cohen has spoken at several conferences and recently chaired the Advanced Learning Institute (ALI) Social Media for Crisis Communications in Government conference in November of 2009.

Shala Ann Byers | has worked for two and a half years as an emergency communications and all-source analyst. She has spent the past year developing a social media reverse mentoring program linking junior staff with senior leadership to facilitate technology and social media learning. Ms. Byers holds a Bachelor's degree from Dartmouth College in Government with a specialty in International Relations.

References

1. *http://www.alexa.com/siteinfo/facebook.com.*
2. *http://www.alexa.com/siteinfo/youtube.com.*
3. *http://www.whitehouse.gov/the_press_office/TransparencyandOpenGovernment*
4. *http://www.openthegovernment.org/otg/OGD.pdf.*
5. *www.bt.cdc.gov/disasters/hurricanes.*
6. *http://www.cio.gov/Documents/Guidelines_for_Secure_Use_Social_Media_v01-0.pdf.*
7. *http://www.fcw.com/Articles/2009/03/25/web-GSA-agreement.aspx.*
8. *http://www.wired.com/dangerroom/2009/01/usaf-blog-respo*

Insider Threat Center at CERT Grows Solutions from Reality-Based Research

by Dawn Cappelli and Andrew P. Moore

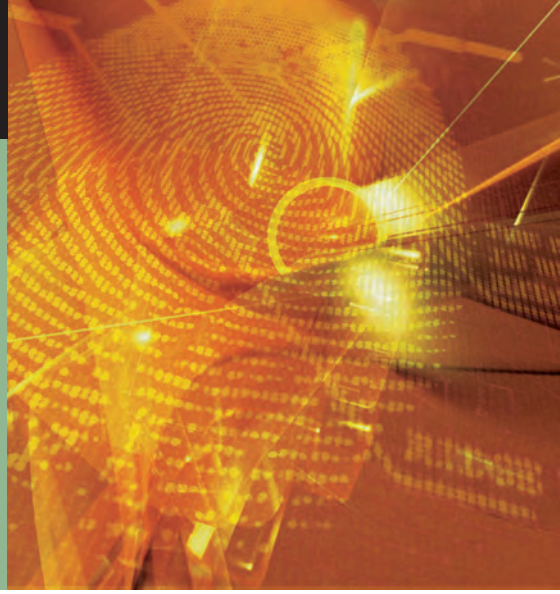

Many organizations have suffered significant losses from insiders with authorized access to protected information assets. Insiders' crimes include theft, sabotage, fraud, and espionage. The Computer Emergency Response Team (CERT), part of the Software Engineering Institute (SEI) at Carnegie Mellon University, began researching this problem in 2001. It has compiled a growing database of more than 300 criminal cases in which current or former employees, contractors, or business partners abused the trust and access associated with their positions. As part of its research, CERT interviewed many of the victim organizations and some perpetrators themselves, complementing a wealth of case data with first-hand insights into the methods and motivations behind these crimes.

This work laid the foundation for the Management and Education of the Risk of Insider Threats (MERIT) project. Under MERIT, CERT researchers collaborated with noted psychologists, the United States Secret Service, the Federal Bureau of Investigation, and the Department of Defense to uncover key technical, social, and organizational patterns of insider behavior. Building on this work, CERT researchers are constructing models of the four main classes of insider crimes: IT sabotage, theft of intellectual property, espionage, and fraud. These models, created using system dynamics techniques, suggest both the evolution of the threat over time and possible mitigation strategies.

Armed with these new insights, the Insider Threat Center at CERT has begun educating organizations on how to detect and manage the problem. It offers its Insider Threat Workshop several times throughout the year. Geared to managers and executives, the two-day workshop addresses technology, organizational culture, policy, procedure, and behavioral issues that influence insider threat. The workshops stress the need to foster cooperation among management, information security, human resources, and IT groups to effectively fight the problem.

CERT has also launched its Insider Threat Vulnerability Assessment program. Spurred by numerous requests from industry and government, these assessments enable organizations to get a better grasp on this complex problem. A CERT project team performs the three-day, on-site assessment, conducting interviews with key organizational personnel. The assessment team explores the organization's technical controls, policies, and [technical and behavioral] practices and then produces a confidential report presenting findings and potential mitigation strategies. The goal is to create a single, actionable framework that engages all stakeholders in the fight against insider threat.

The insider threat team is very excited about the impact it has had on government and industry organizations and their ability to mitigate the risk of insider threat. The workshops and assessments completed to date have proven to be effective tools in raising awareness of the causes, potential indicators, and prevention and detection strategies. CERT now focuses on technical solutions that will enable organizations to use people and technology more effectively."

For more information, please visit *http://www.cert.org/insider_threat/*. ■

About the Authors

Dawn Cappelli | is technical manager of the Threat and Incident Management Group at CERT. She has over 25 years of experience in software engineering, programming, technical project management, information security, and research. She is technical lead of CERT's insider threat research, including the Insider Threat Study conducted jointly by the U.S. Secret Service and CERT.

Andrew P. Moore | is a senior member of the CERT technical staff at the Software Engineering Institute. Moore explores ways to improve the security, survivability, and resiliency of enterprise systems through insider threat and defense modeling, incident processing and analysis, and architecture engineering and analysis. Before joining the SEI in 2000, he worked for the Naval Research Laboratory.

Wikis Within the DoD

by Tzeyoung Max Wu

Wikis within DoD

Web 2.0. Social media is all the hype these days. October 2008 saw the launch of DoDTechipedia, one of the Department of Defense's (DoD) ventures into wikis. Currently, media buzz surrounds the secretive and ambitious A-Space social portal within the Intelligence community. In 2009, the Centers for Disease Control and Prevention (CDC) used social media tools to increase awareness of emerging data about the H1N1 virus. Information was disseminated across YouTube, Facebook and Twitter, where data was quickly assimilated by millions and helped promote health awareness across the public. From proprietary corporate wiki pages to open video blogging forums, we have seen an explosion of all types of social media implementation and usage across sectors both public and private.

Take the case of NASAsphere, a pilot social media study where a social media portal was implemented to test its value to NASA's Jet Propulsion Laboratory (JPL). Within months, the study concluded that participants were sharing information in ways that would have not happened without the tool. Rather than emailing known coworkers for information, NASAsphere users were encouraged to post inquiries for information on the portal. Almost all informational responses to such queries came from users at different NASA centers. [1] By the end of the study, researchers concluded that the portal created a better sense of unity and belonging in NASA participants, despite being separated both physically and organizationally. The site allowed users to openly communicate on a level playing field, removing barriers such as job status and organizational departments. [2]

Wikis

As one popular form of social media, wikis entered mainstream vocabulary with the launch of Wikipedia in 2001. Although the concept of a community-driven encyclopedia had surfaced from time to time for decades, the advent of the Internet finally made it feasible for millions of individual users to freely add and edit content to an open repository of topical articles. By 2008, Wikipedia housed more than 10 million articles, and in 2005, this encyclopedia was pronounced as accurate as the popular Encyclopedia Britannica. [3] Attempting to reap the benefits of seamless community-driven information sharing, corporations and public agencies have since implemented their own proprietary wiki solutions. When wiki solutions work, they provide an enormous amount of value.

Intellipedia, another solution within the government, is a poster-child of wiki success, with core officers earning Homeland Security Awards in 2009. [4] The Intelligence community produces 50,000 reports annually; a majority of them go unread. [5] Amidst data overload, Intellipedia was conceived to promote real-time information sharing internally across the community. It now boasts nearly one million pages and 100,000 users with over 10,000 edits daily. In 2008, following the terrorist bombing of hotels in Mumbai, intelligence analysts convened on a page, created on Intellipedia, to share emerging information and brainstorm ideas. The page received 7,000 views within three days and was integral in the community's analysis of the attack. [6]

DoDTechipedia, itself a relatively new internal wiki solution, run by the Defense Technical Information Center (DTIC), shows much potential for bridging informational silos within DoD. The wiki solution won the 2009 Government Computer News (GCN) Award for agencies. GCN, a news site serving the government market, describes DoDTechipedia as more than a wiki, but rather an entire suite of services spurring collaboration.

Focused DoD Wikis

A set of one or more targeted wiki sites, each effectively addressing the needs of the respective community, can facilitate communication and promote collaboration. Note, 'targeted' is a must for a wiki site. Too broad a scope risks dilution, since at a certain point there is a threshold for the amount of content that must be collected before the site

appears informationally substantial to any specific target community. This is especially true within DoD, where program managers may be more secretive about their research. Thus, the more categories there are, the more content that must be generated to convince communities of its utility. The key is to focus. Of the handful of success factors mentioned by Larry Sanger, one of the founders of Wikipedia, the contribution of a small core group of good people during the early days was key. [7]

A precisely defined target market segment for any DoD wiki site allows for better and speedier marketing to defined communities. With a specified community in mind, the site can be fine-tuned, tailoring everything from look and feel, navigation, editing protocols, registration processes and site promotion to better match the community's needs. For at its core, social media sites, including wikis, have historically been grassroot efforts growing from the bottom up in an organizational hierarchy, with roots deeply tied to their respective user groups. Grassroot efforts survive and mature because they address unmet recognized needs that differ between organizations. As such, participation and content management must remain in the hands of the general contributors so that they are empowered to innovate and run with fresh ideas.

As a grassroots styled site, a wiki needs to become a natural fabric of the community's culture. One of the reasons that Intellipedia worked well was because the custom of social networking, information inquiry and response, and information analysis had already been deeply ingrained into the Intelligence community culture. Part of the challenge for social media sites in DoD will be overcoming a more conservative culture, where informational secrecy has generally been critical to military success and where the sheer size of the organization has necessitated a level of bureaucracy. A successful wiki implementation has to come hand-in-hand with transforming this culture. Facing a similar challenge within the private sector, a human resources firm in Europe devised a comprehensive strategy to build momentum for their internal site. This strategy included employee training, proactive wiki gardening, appointing wiki evangelists and mandating that meetings be recorded and tracked using wiki pages. The latter helped instill into the portal the daily activities of individuals in the firm. [8]

Of course, success cannot happen as a solitary effort. Wikipedia's own success would not have been achievable without the rising popularity of Google's oft-storied search engine. As Google's crawlers started indexing Wikipedia pages, general topical searches on the engine started to return Wikipedia within results. Featuring easy use, open editing, and proven return for efforts, usage of the encyclopedia skyrocketed. Wiki implementations within DoD should be promoted along with complementary solutions and efforts within the organization.

In the end, any wiki implementation must be accompanied with patience and persistence. Intellipedia, itself already springing from an organizational culture deliberately conducive to information gathering, is touted as a success today, but was launched in 2005. The broader the scope of the target communities in the site, the more content that must be generated to reach maturity. Wikipedia, with incredible scope, took many years to garner support from millions of contributors throughout the world. DoD itself has a deeply ingrained conservative culture, with a population of subject matter experts many times smaller. Before the different DoD communities can fully embrace and use wiki sites to their full potential, a degree of culture change will have to occur. One tactic for effective wiki implementation could be to forward social media pilots such as NASAsphere. Pilots can be run for short time periods to measure the site's applicability to the respective needs in the community. Shorter pilots building towards more long-term solutions could be much more cost-effective than a series of failed large-scale efforts.

Security

Of course, information security will remain a key concern, especially with national security at risk. Throughout 2009, DoD wrestled with a balanced social media policy that would allow it to reap benefits, but at an appropriate risk level. There were special concerns about soldiers and other interested parties leaking sensitive operational information on media sites. The US Marine Corps dealt with the security issue by prohibiting all social media use. However, such a policy entirely abdicates the real value that social media can produce. To not fully leverage innovations in technology and media risks DoD falling behind other agencies in the world. In a recent blog post, even Rob Carey, US Navy Chief Information Officer (CIO), said that social media is a resource that DoD should well use to facilitate trust and collaboration. [9] "These tools are fundamental to collaboration. They have the potential to leverage the collective wisdom of this 750,000+ member Department," said Carey.

Security risks are real, but can be strategically mitigated to a certain degree *via* a smart architecture and set of policies. One interesting solution described on the Armed Forces Communications and Electronics Associate (AFCEA) Web site proposes setting up dedicated Internet services for all staff. [10] Internet services centralized in this way allow administrators and automated tools to better scan information posted to the Internet and catch security data leaks more effectively. This could be a broader social computing solution for computer use on the global information grid (GIG) in general, where bare-boned computer terminals plug onto resources served and managed on the GIG, providing a set of virtual desktops to users wherever they can plug into the GIG.

Any technical solution must be coupled with DoD guiding policies as well as real culture change. In September 2009, the Federal CIO Council issued official guidelines for Secure Use of Social Media by Federal Departments and Agencies. [11] The very first risk mitigation step suggested was the need for a government-wide policy for social media that would address policy controls, acquisition controls, training controls, and host and network controls. The guidelines define four types of information traffic that must be managed: inward sharing, outward sharing, inbound sharing, and outbound sharing. Each of these four types of information flow come with unique risks and mitigation approaches. From a cultural perspective, DoD users should be trained with a practical sense of caution when utilizing social media systems.

Wikis within DoD will require a fair amount of monitoring, both from a content perspective as well as in network security and information assurance. A cultural shift toward data sharing and collaboration should also be tempered with an appropriate culture of caution and sensibility within the user community. This is quite achievable, of course, and will be important in the ongoing evolution of DoD to accomplish its missions in the hastening change of technology. Collaboration will accelerate the pace of innovative problem resolution within DoD. ■

About the Author

Tzeyoung Max Wu | was a DoDTechipedia content manager, creating and editing material in IA, information warfare, and networking technology areas. His experiences in information technology security have included: administering and configuring servers and network devices within organizations; designing secure architecture for enterprise systems; and configuring access control lists, profiles, and border controls for network applications.

Mr. Wu received his Bachelor's degree in computer science from New York University, holds an MBA at the University of Chicago Booth School of Business, and earned a Master's degree in IT from Virginia Tech.

References

1. Jackson, Joab. NASA program proves the benefits of social networking. Government Computer News. 2009. *http://www.gcn.com/Articles/2009/11/30/A-Space-side-NASA-social-networking.aspx* (accessed 01/02/2010).
2. Merryman, Celeste. Findings from the NASAsphere Pilot. Jet Propulsion Laboratory, California Institute of Technology Knowledge Architecture and Technology Task. (Pilot team: Merryman, Celeste; Hughes, Dougals). California Institute of Technology. 2008. *http://www.scribd.com/doc/12759868/NASAsphere-Pilot-Report-2008-Public* (accessed 01/02/2010).
3. Terdiman, Daniel. Wikipedia hits 10 million total articles. CNET. 2008. *http://news.cnet.com/8301-13772_3-9905726-52.html* (accessed 01/02/2010).
4. Intellipedia Gurus Win 2009 Homeland Security Medal. CIA website. *https://www.cia.gov/news-information/featured-story-archive/intellipedia-homeland-security-medal.html* (accessed 01/02/2010).
5. Thompson, Clive. Open-Source Spying. The New York Times. 2006. *http://www.nytimes.com/2006/12/03/magazine/03intelligence.html* (accessed 01/02/2010).
6. Intellipedia Gurus Win 2009 Homeland Security Medal. CIA website. *https://www.cia.gov/news-information/featured-story-archive/intellipedia-homeland-security-medal.html* (accessed 01/02/2010).
7. The Early History of Nupedia and Wikipedia, Part II. Slashdot. *http://features.slashdot.org/article.pl?sid=05/04/19/1746205* (accessed 01/02/2010).
8. Roberts, Bill. How to Marshal wikis: some human resource professionals are using wikis to communicate, collaborate. HR Magazine. 2008. *http://findarticles.com/p/articles/mi_m3495/is_12_53/ai_n31159337/pg_2/?tag=content;col1* (accessed 01/02/2010).
9. Carey, Rob. Embracing Social Networking Tools. Department of the Navy CIO. 2010 *http://www.doncio.navy.mil/Blog.aspx?ID=891* (accessed 2/3/2010).
10. Strassman, Paul A. Social (Network Security. Signal Online. 2010 *http://www.afcea.org/signal/articles/templates/Signal_Article_Template.asp?articleid=2163&zoneid=284* (accessed 2/1/2010).
11. Guidelines for Secure use of Social Media by Federal Departments and Agencies, v1.0 *http://www.doncio.navy.mil/Download.aspx?AttachID=1105* (accessed 2/3/2010).

Penn State Industry Day Conference

by Rich Coulter

The Networking and Security Research Center (NSRC) at the Pennsylvania State University held its annual Industry Day from 13 to 14 October 2009 at the University Park campus in State College, Pennsylvania. The NSRC provides a research and education community at Penn State for professors, students, and industry collaborators interested in networking and security. Industry Day is an opportunity for partners and other interested industry members to learn about research over the past year and ongoing developments.

Dr. Frank Siebenlist and Robin Burk delivered keynote addresses. Dr. Seibenlist is a senior security architect at the Mathematics and Computer Science Division at the Department of Energy Argonne National Laboratory and a Fellow at the Computation Institute of the University of Chicago. Ms. Burk currently manages the basic research thrust in cognitive, information, and network science for the Defense Threat Reduction Agency .

Dr. Tom La Porta, NSRC Director, noted that two NSRC faculty members received National Science Foundation (NSF) Presidential Early Career Awards for Scientists and Engineers in 2009. Only 25 of these prestigious awards are presented each year, so it was a truly unique event for two faculty from the same university to receive them. Dr. Adam Smith was recognized for his work in data access and privacy, and Dr. Sean Hallgren was awarded for developments in quantum computation.

Dr. Patrick McDaniel, co-director of the Systems and Internet Infrastructure Security (SIIS) laboratory presented analysis of several networked devices intended to monitor and control electrical power usage for a "smart grid." The SIIS lab discovered vulnerabilities that could be exploited to overload generation plants, deny power to critical customers, or obfuscate power usage. Dr. McDaniel is also exploring attack causality in Internet-connected cellular networks with the goal to understand and protect against evolving threats in cellular phone systems. Other ongoing projects in the SIIS lab include Telecommunications Security; Voting Systems Integrity; and security of systems, virtual machines (VM), and storage.

Each graduate student in the NSRC also presented posters summarizing their research. Their research focused on networking (security, fault isolation, coding, efficiency, encryption), mobile devices (device security, network threats), and systems (VM security policy, software theft detection).

Other affiliated Penn State resources for industry were highlighted. The Applied Research Laboratory (ARL) is a DoD-designated U.S. Navy University Affiliated Research Center that maintains a long-term strategic relationship with the Navy and supports the other services as well as industry. ARL also provides facilities for conducting classified work in conjunction with the NSRC. The Industrial Research Office (IRO) focuses on uncovering researchers in all Penn State colleges and departments to meet industry needs. IRO facilitates industry partnerships with the NSRC and other research centers at Penn State.

Briefings can be found at *http://nsrc.cse.psu.edu/id09.html*. More information on ARL and the IRO can be found at *http://www.arl.psu.edu/* and *http://www.research.psu.edu/iro/index.asp*, respectively. ■

About the Author

Richard Coulter | currently provides remote systems engineering and project management support on various projects, and works to establish relationships between IATAC and Penn State, especially in support of the Administration's Cybersecurity Initiative. Previously, Mr. Coulter performed hardware and embedded design, reverse engineering, and data analysis in support of law enforcement forensic and operational missions, where he served as deputy program manager. Mr. Coulter received a Bachelor's degree in electrical engineering from the Pennsylvania State University.

Vulnerability Assessment Processes Within DoD

The Problem

Protecting critical infrastructure and the Global Information Grid continues to be a valuable, yet time-consuming and expensive effort within the Department of Defense (DoD). Initiatives and compliance requirements including Federal Information Security Management Act, the Federal Desktop Core Configuration, Computer Network Defense Service Provider compliance efforts, mandates from the Joint Task Force – Global Network Operations (JTF-GNO) and the Defense Information Systems Agency (DISA), and general due diligence to protect the technology and data that keeps the U.S. military operational are iterative, redundant, and in many cases, based on manual processes.

Configuration management, patch management, and vulnerability and risk management are all predicated upon processes that are cyclical and typically involve hands-on efforts by system or network administrators. They may also require compliance reviews from information assurance divisions, testing from vendors and system managers, approval from configuration control boards, and ultimate acceptance from the Designated Accrediting Authority for the organization, system, or enclave. In many cases, the process of assessing compliance and validating appropriate configuration, and more importantly, identifying weaknesses and vulnerabilities within established configurations, is accomplished by performing vulnerability assessments.

Vulnerability assessment processes in many organizations are ad-hoc, non-standardized, and incomplete. They rely on commercially developed tools as well as DoD-provided tools and in-house solutions to determine patch levels, user settings, open ports, operating system configurations, and other system (mis)configurations. Unfortunately, no one vulnerability assessment solution is comprehensive enough to cover all niches and corners of the DoD infrastructure. Because of this problem, technologists and oversight organizations are required to use multiple vulnerability assessment tools to help ensure that all bases are covered. Some assessment tools are proficient at scanning Microsoft Windows; some are good for UNIX-based operating systems; some excel in evaluating Web applications; and others do device discovery very well. The shape and composition of the environment often dictates what tools need to be used to manage compliance and ensure secure configuration whenever possible.

Having to rely on multiple vulnerability assessment solutions means that technologists and oversight personnel are reduced to seeing vulnerability and configuration data in many disparate, non-standard views. This can make managing and tracking efforts to meet compliance goals and secure the infrastructure exceptionally difficult, because no standardization exists across the entire enterprise. This problem is compounded by employee or contractor turnover, the volatility in technical or mobile environments, and the various skill levels of personnel working to manage the infrastructure. It is also exaggerated by the fact that vulnerability assessments and compliance scans play such a big role in major DoD programs and mandates that include the information assurance vulnerability management process, certification and accreditation, computer network defense, information operations condition, and JTF-GNO mandates.

Recommended Solutions

The first place to begin addressing compliance and configuration management issues is to have an overarching configuration management plan. It is crucial to have a healthy cross-section of the technologists within the organization designated as members of a configuration control board (CCB) that is strictly governed by documented configuration management processes and procedures. As part of that configuration management plan, however, there also need to be specific guidelines and instructions on how to perform vulnerability assessments within the organization to ensure

appropriate configuration and validate the mandates of the DoD as interpreted and implemented by the CCB. This vulnerability assessment process should be created and maintained by the personnel responsible for implementation of the technology as well as those areas of the organization that are responsible for oversight and compliance reporting. The primary goal of the plan should be to standardize the process, make it repeatable, and enforce it for all vulnerability assessment activities.

A vulnerability assessment manual for an organization should address and define procedures for several key components of the vulnerability assessment process. These areas include:

- **Approved vulnerability assessment tools list**—It is important to ensure that senior management (the chief information officer [CIO] or chief information security officer [CISO]) acknowledges what tools are permitted to be used within the network or enclave. To this end, a formal memo drafted by the CIO/CISO should specifically designate vulnerability assessment tools that are approved for use and prohibit the use of any tools not explicitly allowed. This will help ensure that untested, unknown vulnerability assessment tools do not adversely impact operations of the network or enclave and ultimately thwart the mission of the organization.
- **Specific attributes and definition of each tool**—Each approved tool has information that needs to be maintained and remains relevant for the life of the tool. Support information, update processes, training materials, known issues with the tool, the types of targets the tool is capable of assessing—these are the kinds of things that need to be recorded and kept up to date to ensure that anyone required to perform a vulnerability assessment has the appropriate information to do so effectively.
- **Process for coordinating and approving vulnerability assessments**—Sufficiently defining this step is one of the most important goals of any vulnerability assessment manual. A standardized test matrix should be developed and used to define and coordinate any vulnerability assessment activities. The test matrix should include information such as the targets, tools to be used, ports to be scanned, scan policy to be used, scan throttling information, points of contact, and date and time of the scan. The test matrix should be used to coordinate with components that may be impacted by the assessment—the system manager, program manager, network monitors, and even users.
- **Process for consolidating, distributing, and storing assessment results**—The point of a vulnerability assessment manual is to standardize processes and make them repeatable. As such, this is also a very important part of the process. The plan should outline acceptable formats for vulnerability assessment results. If results from disparate tools are aggregated or consolidated in any way, the process used to do that should be outlined. Where and how the vulnerability and configuration information is stored should also be specifically outlined. Emerging technology has been developed to facilitate this process and help bridge the reporting gap between separate vulnerability assessment tools.
- **Troubleshooting vulnerability assessments and the correlation to incident response**—Troubleshooting vulnerability assessment tools are also paramount to standardization. If tools are not used or are not functioning correctly, results can be skewed and the configuration and security posture of the targets scanned may not be accurate. It is also important to remember

(especially for legacy systems), that there is potential to bring down production systems if they are targeted intentionally or unintentionally. The vulnerability assessment process should identify incident response procedures in the event that an assessment causes an outage or adverse reactions by the targets being scanned.

Incorporating these types of guidelines and parameters into a vulnerability assessment plan is vital. Without standardization and appropriate training to perform vulnerability assessments, it is easy to have vulnerabilities or misconfiguration missed—ultimately resulting in a false sense of security for the organization and greater risk to the mission and the DoD.

Also, don't be afraid to leverage virtualization. Virtualization can be a great tool in the vulnerability assessment space—especially in environments with legacy systems and antiquated technology. Using virtualization to take an exact copy of a production server or application allows for extensive vulnerability assessment that may otherwise not be possible.

Options

Establishing (and following) a vulnerability assessment manual as part of a bigger configuration management plan is not difficult, and it is not exceptionally time consuming. In fact, implementing a standard approach to vulnerability assessment activities can ultimately save a lot of time and effort by streamlining the process and making sure that all relevant vulnerability assessment information can be found in one easy-to-use location.

However, if vulnerability assessments are conducted at recommended (not just required) intervals, agencies within the DoD may find that adhering to rigorous vulnerability assessment processes can be expensive and time consuming—especially in larger, more distributed environments. It is for this reason that many organizations merely do what is specifically required by JTF-GNO or DISA or any other oversight organization with the ability to push down DoD requirements.

Performing the scans is not generally the difficult or time-consuming part of the process; it is interpreting, processing, and putting to work the volumes of information that the vulnerability assessment tools return—especially given the points discussed above. Using only one or two vulnerability assessment solutions for most organizations is insufficient, especially within the DoD. So consolidating, aggregating, and presenting the results of disparate vulnerability assessment scans is generally the most resource-intensive part of the process.

Organizations have two options. The first is to rely on the native outputs of the various vulnerability tools themselves. This could be flat text files, XML files, HTML files, PDFs, or Microsoft Word documents. For some tools, it could even mean having to rely on the console of the vulnerability assessment tool itself instead of a report. In this scenario, presenting findings in terms of high, medium, and low risk is disjointed and subject to error. It also makes remediation efforts difficult for system and network administrators because they have to rely on so many different forms of information from the various assessment tools that do not look similar and do not always present the most useful information.

The second option includes processes of trying to manually consolidate the data to put it into a more meaningful/useful format that facilitates the efforts of administrators and makes tracking progress a bit easier. The problem with this scenario is that it is full of manual copying and pasting, parsing, or scripting that is not vetted or standardized, and it remains exceptionally time consuming.

Great strides have been made to facilitate resolution to this problem. New, emerging technologies attack this problem head-on by providing the capability to consolidate, aggregate, and re-present vulnerability information in a truly meaningful fashion. The process of consolidating vulnerability data for system administrators no longer takes days and hours; with the right solution, it can take only minutes.

Conclusion

One of the most important pieces of the configuration management process is inspection and validation through vulnerability and configuration assessments. These processes can be time consuming; however, their value is obvious, and they also play fundamental roles in other major programs and initiatives implemented by the DoD. It is critical to have standardized processes when it comes to vulnerability assessments because when ad-hoc processes fail, and they do too often, it is difficult to trust the outcome of those assessments, and making decisions based upon misinformation can be devastating.

Armed with a thorough and well-implemented vulnerability assessment plan and with new technology that allows system and network administrators to focus more on resolving vulnerabilities and misconfiguration and less on combing through volumes of data for useful information, maintaining compliance with fewer resources becomes reality. ■

About the Author

Chris Merritt | is the president and CEO of Prolific Solutions, LLC (*www.prolific-solutions.net*) and has been consulting for the DoD for over seven years. He is the author of proVM Auditor (*www.provmauditor.com*), a vulnerability assessment aggregation and compilation tool, and holds a number of information security certifications, including CISSP and CISA. He earned his Master's degree in information assurance from Norwich University in 2007.

Dr. Peng Liu

by Angela Orebaugh

This article continues our profile series of members of the Information Assurance Technology Analysis Center (IATAC) Subject Matter Expert (SME) program. The SME profiled in this article is Dr. Peng Liu from Pennsylvania State University.

Dr. Peng Liu is an Associate Professor in the College of Information Sciences and Technology (IST). He is also a member of the graduate faculty for the Department of Computer Science and Engineering and affiliate associate professor for the Department of Supply Chain and Information Systems (SC&IS) in the Smeal College of Business. In addition, Dr. Liu is the Director of the Cyber Security Lab and Director of the LIONS Center. His research interests include survivable systems, systems security, information security, network security, privacy, identity theft, cyber infrastructures, and electronic health. [1]

Dr. Liu won a $6.25M grant from the Army Research Office in July 2009 to study cyber situation awareness (CSA). He and his team received a Multidisciplinary University Research Initiative Award (MURI) for his project, "Computer-aided Human-centric Cyber Situation Awareness." They plan to use the grant funding to further the research on cyber awareness and how it can be used to improve cyber defense. Research goals include developing tools that will help bridge the gap between analysts' capabilities and existing CSA software and hardware. The objective of this effort is to develop an integrated end-to-end (spanning the whole 'life cycle') CSA solution to fill the gap between machine information processing and analysts' mental processes. The scope of this effort is to develop new capabilities for computer-aided human-centric CSA. The solution adds the new algorithms and techniques that are needed for the machine situational awareness (SA) system to work in concert with the human SA system. It integrates the human cognition aspects and the computer algorithm aspects of cyber SA. The solution also integrates situation recognition, impact assessment, causality analysis, trend analysis, and assessment of system assurance. The team will develop prototype capabilities in each year of the project that build on prior years' capabilities, with the goal of having a testable, executable prototype at each stage of the project.

Dr. Liu was also one of three researchers who received more than $1M funded by the American Recovery and Reinvestment Act of 2009. His project—Collaborative Research: Towards Self-Protecting Data Centers: A Systematic Approach—is aimed at safeguarding business applications and infrastructure from cyber threats. The research team seeks to improve security consolidation to meet the top two requirements for modern data centers—business continuity and information security. The team will take a systematic approach that leverages the emerging virtual machine technologies to consolidate four areas of systems security research: microscopic intrusion analysis and detection; redundancy; automatic response; and diversity-driven protection. Broader impacts for this research include a significant advancement in reducing risks to business applications and information systems, increasing business continuity, and delivering data assurance in the presence of severe cyber attacks. Liu will co-lead this project, which will further the team's previous research on cyber awareness and how it can be used to improve cyber defense.

Dr. Liu organizes and presents at several conferences in information security. A few examples include: Securecomm 2009 (general chair); Inscrypt 2008 (both Program Co-Chair and keynote speaker); and AsiaCSS 2010 (Program Co-Chair). ■

References

1. *http://ist.psu.edu/s2/pliu*

Eight Steps to Holistic Database Security

by Dr. Ron Ben Natan

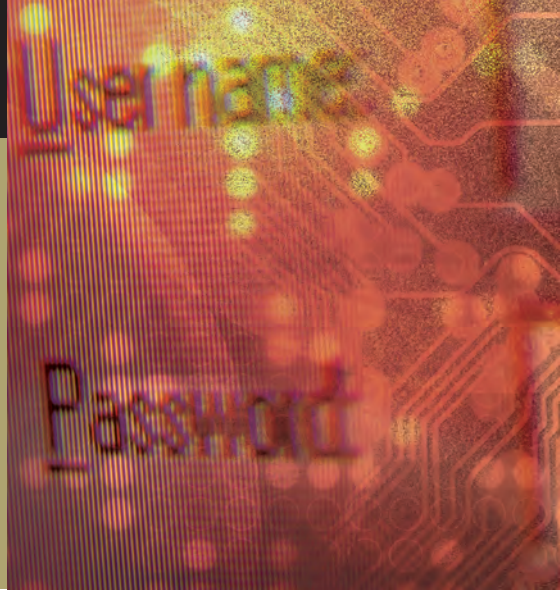

Financially motivated attacks, malfeasance by insiders, and regulatory requirements such as the Federal Information Security Management Act-mandated National Institute of Standards and Technology (NIST) 800-53 standard are driving government organizations to find new ways to secure their data.

Most of the world's sensitive data is stored in commercial database systems such as Oracle, Microsoft SQL Server, IBM DB2, and Sybase—making databases an increasingly favorite target for criminals. This may explain why external attacks such as SQL injection jumped 134% in 2008, increasing from an average of a few thousand per day to several hundred thousand per day, according to a report recently published by IBM. [1]

To make matters worse, according to a study published in February 2009 by the Independent Oracle Users Group (IOUG), nearly half of all Oracle users are at least two or more patch cycles behind in their database patching. [2] In addition, 74% of all Web application vulnerabilities disclosed in 2008 did not even have an available patch by the end of 2008, according to IBM. [3]

Whereas most attention has previously been focused on securing network perimeters and client systems (*e.g.,* firewalls, IDS/IPS, and anti-virus), we are now entering a new phase where information security professionals are now being tasked with ensuring that critical databases are secure from breaches and unauthorized changes.

Here are eight essential best practices that provide a holistic approach to both safeguarding databases and achieving compliance with key regulations and standards such as NIST 800-53 and Defense Information System Agency Security Technical Implementation Guides as well as the Sarbanes-Oxley Act (SOX), Payment Card Industry Data Security Standard (PCI-DSS), and data protection laws:

- **Discovery**—You cannot secure what you do not know. You need to have a good mapping of your sensitive assets—both of your database instances and your sensitive data inside the databases. Plus, you should automate the discovery process because the location of sensitive data is constantly changing due to changes such as new or modified applications and mergers and acquisitions. In an interesting twist, some discovery tools can also find malware placed in your database as a result of SQL injection attacks. In addition to exposing confidential information, SQL injection vulnerabilities allow attackers to embed other attacks inside the database that can then be used against visitors to the Web site.

SQL injection jumped 134% in 2008, increasing from an average of a few thousand per day to several hundred thousand per day.

- **Vulnerability and Configuration Assessment**—You need to assess the configuration of your databases to ensure they do not have security holes. This includes verifying both the way the database is installed on the operating system (*e.g.,* checking file privileges for database configuration files and executables) and configuration options within the database itself (such as how many failed logins will result in a locked account, or which privileges have been assigned to critical tables). Plus, you need to verify that you are not running database versions with known vulnerabilities. Traditional network vulnerability scanners were not designed for this because they do not have embedded knowledge about database structures and expected

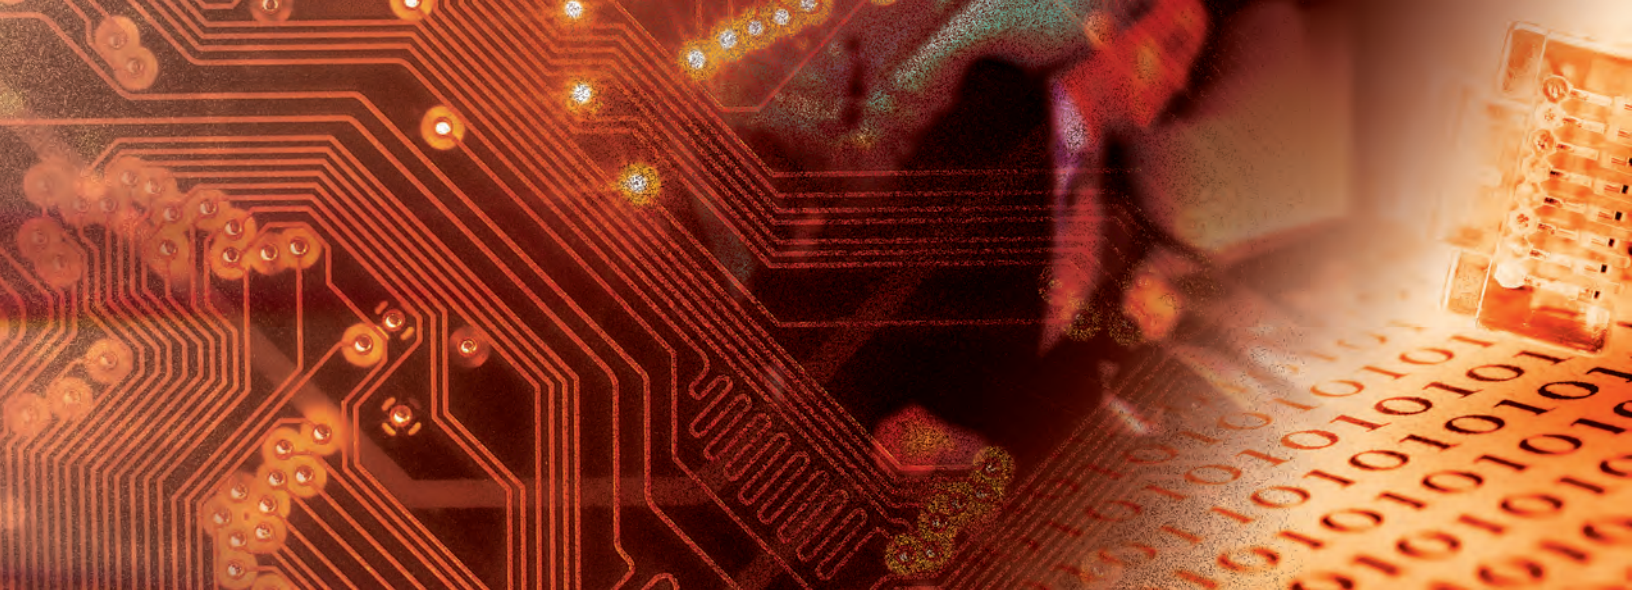

behavior, nor can they issue SQL queries (*via* credentialed access to the database) in order to reveal database configuration information.

- **Hardening**—The result of a vulnerability assessment is often a set of specific recommendations. This is the first step in hardening the database. Other elements of hardening involve removing all functions and options that you do not use.
- **Change Auditing**—Once you have created a hardened configuration, you must continually track it to ensure that you do not digress from your "gold" (secure) configuration. You can do this with change auditing tools that compare snapshots of the configurations (at both the operating system level and at the database level) and immediately alert you whenever a change is made that could affect the security of the database.
- **Database Activity Monitoring (DAM)**—Real-time monitoring of database activity is key to limiting your exposure by immediately detecting intrusions and misuse. For example, DAM can alert on unusual access patterns indicating a SQL injection attack, unauthorized changes to financial data, elevation of account privileges, and configuration changes executed *via* SQL commands. Monitoring privileged users is also a requirement for data governance regulations such as SOX and data privacy regulations such as PCI-DSS. It is also important for detecting intrusions because attacks will frequently result in the attacker gaining privileged user access (such as *via* credentials owned by your business applications). DAM is also an essential element of vulnerability assessment because it allows you to go beyond traditional static assessments to include dynamic assessments of "behavioral vulnerabilities" such as multiple users sharing privileged credentials or an excessive number of failed database logins. Finally, some DAM technologies offer application-layer monitoring, allowing you to detect fraud conducted through multi-tier applications such as PeopleSoft, SAP, and Oracle e-Business Suite, rather than through direct connections to the database.
- **Auditing**—Secure, non-repudiable audit trails must be generated and maintained for any database activities that impact security posture, data integrity, or viewing sensitive data. In addition to being a key compliance requirement, having granular audit trails is also important for forensic investigations. Most organizations currently employ some form of manual auditing, utilizing traditional native database logging capabilities. However, these approaches are often found to be lacking because of their complexity and high operational costs due to manual efforts. Other disadvantages include high performance overhead, lack of separation of duties (because database administrators can easily tamper with the contents of database logs, thereby affecting non-repudiation) and the need to purchase and manage large amounts of storage capacity to handle massive amounts of unfiltered transaction information. Fortunately, a new class of DAM solutions are now available that provide granular, database management system (DBMS)-independent auditing with minimal impact on performance, while reducing operational costs through automation, centralized cross DBMS policies and audit repositories, filtering, and compression.
- **Authentication, Access Control, and Entitlement Management**—Not all data and not all users are created equally. You must authenticate users, ensure full accountability per user, and manage privileges to limit access to data. And you should enforce these privileges—even for the most

privileged database users. You also need to periodically review entitlement reports (also called User Right Attestation reports) as part of a formal audit process.

- **Encryption**—Use encryption to render sensitive data unreadable, so that an attacker cannot gain unauthorized access to data from outside the database. This includes both encryption of data-in-transit, so that an attacker cannot eavesdrop at the networking layer and gain access to the data when it is sent to the database client, as well as encryption of data-at-rest, so that an attacker cannot extract the data even with access to the media files.

A holistic database security approach is needed to protect against cyberattacks, breaches, fraud, and insider threats. Additionally, such a strategy helps federal agencies and contractors meet NIST 800-53 and comply with the OMB M-06-16 directive, Protection of Sensitive Agency Information, in order to secure personally identifiable information and other sensitive data such as financial data and classified information. ■

About the Author

Dr. Ron Ben Natan | chief technology officer for Guardium, the database security company, has more than 20 years of experience developing enterprise applications and security technology. Guardium, an IBM Company, delivers a scalable platform that prevents information leaks from the data center and ensures the integrity of enterprise data. The company's enterprise security platform is now installed in more than 450 data centers worldwide, including top government agencies. Dr. Natan has authored 12 technical books, including *HOWTO Secure and Audit Oracle 10g and 11g* (© 2009 by Taylor and Francis Group, LLC) and *Implementing Database Security and Auditing* (© 2005, Elsevier, Inc.), the standard texts in the field.

References

1. IBM Global Technology Services, "IBM Internet Security Systems X-Force® 2008 Trend & Risk Report," January 2009.
2. IOUG, "Security Patching Practices by Oracle Users," February 2009.
3. Ibid.

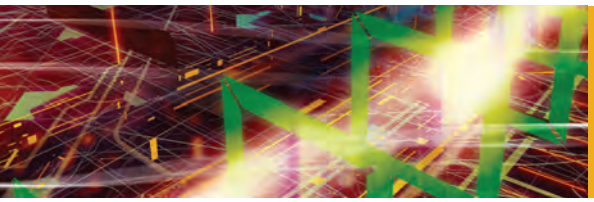

Letter to the Editor

Q ***There are a lot of information assurance conferences, forums, and seminars available to the IA community, and the IAnewsletter focuses on several each year. What is the most important IA conference IATAC takes part in annually?***

A A critical aspect of sharing information assurance (IA) related information is attending events where solutions for pressing IA problems can be discussed. These events also help the IA community learn about the resources available to them and some of the cutting-edge developments in the IA field. IATAC attends, exhibits, and presents at several conferences a year to take part in critical IA discussions, and to promote outreach and awareness for the free products and services we offer. The biggest conference we attend each year is the **Information Assurance Symposium (IAS)**, hosted by the National Security Agency, Defense Information Systems Agency, and US Strategic Command.

This year's conference took place in Nashville, TN, February 2-4, bringing together over 2,000 attendees from all three of IATAC's target communities: government, industry, and academia. Attendees had the opportunity to participate in one of four tracks. The *Protect* track focused on discovering ways to improve information security and harden networks. The *Defend* track looked at how cyber warriors can detect, diagnose, and respond to security threats effectively. The *Survive* track featured sessions on sustaining mission essential functionalities during network attacks. Finally, the *Making it all Happen* track analyzed how to staff, equip, train, and certify the cyber warrior.

IAS stressed the importance of true collaboration and the need to achieve information superiority, and it provided the IA community with networking opportunities essential to achieving these goals. IATAC was glad to take part in IAS this year, and we look forward to participating again next year. ■

Public/Private Partnership Becoming a Necessity

by Allan Carey

Governments have long dealt with espionage and attempts to exfiltrate state secrets and intellectual property. The interconnected world of computing systems has split our efforts to detect and thwart such attempts between the physical and logical worlds. The term advanced persistent threat (APT) has had relevancy in the information assurance world, which started in the US Air Force around 2006. However, beyond government and the defense industrial base, no one in the private sector had really heard or cared about APT.

Until now…Why? Google *vs.* China catapulted APT into the mass media spotlight for better or worse. [1] Back in July 2009, Richard Bejtlich ran a Google search on "advanced persistent threat" prior to an Institute for Applied Network Security briefing which yielded 34 unique hits. [2] As of 16 January 2010, the same search returned 169 hits. During the week of 25 January 2010, The Christian Science Monitor reported about stolen bid data from three major energy companies with traces back to China. [3] And Mandiant, a specialized consulting firm, released its first M-Trends Report which highlighted the types of attacks they have investigated including ones perpetrated by the APT. [4]

Let's start with the negative part of this attention. APT has just made the buzzword bingo chart of marketing professionals targeting our industry. The term will be misrepresented, misused and basically abused to promote/sell products and services with the promise of solving this problem. For the misguided, their attention and resources will be directed away from solving their real information assurance problems. For the well informed, they should see right through the APT elixir.

On the positive side, senior security leaders are now more aware of this threat vector, even though they may not have the budget or resources to do something about it. As a result, organizations are getting engaged in the conversation and looking for ways to collaborate and share information. Changing the way in which we interact and exchange best practices must occur, particularly around this topic, because our advanced persistent adversaries, are incredibly organized and well funded. They are sharing best practices and techniques; as a profession, we must do the same because continuing to fight the battle in silo efforts is not a sustainable strategy.

One promising example of public/private partnership is the impending Google and the National Security Agency relationship. This action is a step in the right direction for sharing defensive techniques and enabling another organization to better defend itself. Another example is the National Security Telecommunications Advisory Committee Network Security Information Exchanges, which I believe will see increased participation from industry in light of the recent developments. Other groups/relationships are forming behind closed doors, but the motivation and business drivers are strong enough to hopefully change the paradigm between public/private partnership and information sharing overall. ■

References

1. *http://googleblog.blogspot.com/2010/01/new-approach-to-china.html.*
2. *www.taosecurity.com.*
3. *http://www.csmonitor.com/Commentary/editors-blog/2010/0126/Why-the-China-virus-hack-at-US-energy-companies-is-worrisome.*
4. *http://www.mandiant.com/news_events/article/mandiant_releases_first_annual_m-trends_report_at_u.s._department_of_d*

Apples & Oranges: Operating and Defending the Global Information Grid

by Dr. Robert F. Mills, Major Michael B. Birdwell, and Major Kevin R. Beeker

Cyberspace is a contested, warfighting domain, but we're not really treating it as such, partly because our language and doctrine have not matured to the point that allows us to do so. One reflection of our immature language is our inability to clearly differentiate the concepts of *network operations (NETOPS)* and *computer network defense (CND)*. This creates confusion about the roles and responsibilities for provisioning, sustaining, and defending the network—much less actually using it. In this article, we resolve this confusion by highlighting the differences among maintenance, defense, and mission assurance activities. Only by separating these activities can we more effectively organize, train, and equip people to perform those tasks. We also describe how the mission assurance aspect of NETOPS can better be viewed as a force protection issue, thereby highlighting the importance of the unit commander in the cyberspace puzzle.

Culture Change

There has been much talk about changing our cyber culture—specifically with respect to how we use cyberspace. General Kevin J. Chilton, the Commander of US Strategic Command (USSTRATCOM), hosted a Cyberspace Symposium in April 2009. In his opening remarks, he labeled cyberspace operations as commanders' business and described a shift in culture that must occur for the United States to be effective in this domain: "We must think about this domain and the tools in this domain and the readiness of this domain as commanders, as essential to successful operations." General Chilton calls every Soldier, Sailor, Airman, Marine, DoD civilian, and contractor to arms, saying, "They are part of the front line of defense and in fact they're engaged in cyber operations that matter every day, whether they know it or not." He compares operations in the domain to "the guards who guard your bases, who stand there at the gate and make sure only the right people come in and keep the wrong people out—that's everybody who has a computer on their desk in these domains today." [1]

Similarly, Air Force Chief of Staff General Norton A. Schwartz sent an e-mail to every member of the Air Force entitled *Cyberspace Operations Culture Change* on May 27th, 2009. In this e-mail he wrote, "Compliance with time critical software updates will gain new emphasis and commanders will be held accountable…. Our Air Force must move to a system of tight network control, personal responsibility, and accountability as we execute our global mission on behalf of our Nation." [2] General Schwartz made it clear that all Air Force members operate in cyberspace and echoed General Chilton's comments emphasizing commander involvement and responsibility for cyberspace operations.

Our leaders are making some very interesting points here. We are all on the front line of defense and are involved in cyber operations every day. General Chilton's analogy of the gate guard who "keeps the wrong people out" is noteworthy, but his use of the word 'defense' is misleading—he's really talking about 'security and force protection.' But he's not the only one who falls into this trap—our doctrine is just as confusing.

NETOPS and Network Defense

This is how the DoD Dictionary defines NETOPS and CND:

- **NETOPS**—"activities conducted to operate and defend the Global Information Grid."
- **CND**—"actions taken to protect, monitor, analyze, detect, and respond to unauthorized activity within DoD information systems and computer networks." [3]

Figure 1 illustrates the NETOPS continuum, and demonstrates the difficulty in distinguishing between the two disparate functions of maintenance and defense.

Effective CND uses a defense-in-depth strategy and employs intelligence, counterintelligence, law enforcement, and other military capabilities as required. However, the CND culture is

Figure 1 NETOPS and CND Continuum

largely one of information assurance (*e.g.,* confidentiality, integrity, and availability), system interoperability, and operations and maintenance (O&M). Many of the things that we routinely call 'cyberspace defense' in cyberspace are really just O&M activities—such as setting firewall rules, patching servers and workstations, monitoring audit logs, and troubleshooting circuit problems.

We talk about vulnerabilities and the thousands of 'cyber attacks' against our networks every day, but we do not treat cyberspace operations like those conducted in other domains. Server availability and communications circuit status are represented as green, yellow, and red lights on a stop-light chart, with an objective being 'all green.' And yet, when a system or circuit is reported as yellow or red, we rarely understand what the true operational impact is in a timely manner. Furthermore, thousands of systems administrators routinely count and scan computers to ensure that their software and operating system patches are current. The objective is 100% compliance, but even if we could achieve that, this is a *maintenance* activity. (Indeed, do we even really know how many computers we have, let alone how many are compliant?) This is no more a defensive activity than counting all the rifles in an infantry company and inspecting them to ensure that they are properly cleaned and in working order.

Our current NETOPS/CND mindset is intentionally focused inward, with emphasis on ensuring that friendly forces have freedom of action within and through cyberspace. Contrast this with a traditional warfighting mentality in which we study an adversary's potential courses of action, develop and refine operational plans to meet national and military objectives, parry thrusts, and launch counter attacks. While we do worry about internal issues such as security, force protection, logistics, and sustainment, our focus remains outward on the adversary. Granted, terms such as 'inward' and 'outward' mean different things when discussing cyberspace (because geographic boundaries are somewhat irrelevant), but we generally use these terms to refer to friendly forces and adversaries, respectively.

Our intent is not to diminish the importance of NETOPS activities—these activities are critical to our ability to operate in and through cyberspace. But they are not defensive activities—at least not in the classical understanding of the concept. Turning to Carl von Clausewitz, we see a much different concept of defense than is currently applied to cyberspace:

> *What is the concept of defense? The parrying of a blow. What is its characteristic feature? Awaiting the blow. It is this feature which turns any action into a defensive one; it is the only test by which defense can be distinguished from attack in war. Pure defense, however, would be completely contrary to the idea of war, since it would mean that only one side was waging it.… But if we are really waging war, we must return the enemy's blows; and these offensive acts in a defensive war come under the heading of 'defense' –in other words, our offensive takes place within our own positions or theater of operations. Thus, a defensive campaign can be fought with offensive battles, and in a defensive battle, we can employ our divisions offensively. Even in a defensive position awaiting the enemy assault, our bullets take the offensive. So the defensive form of war is not a simple shield, but a shield made up of well-directed blows. [4]*

Similarly, Army Field Manual 3-0, Operations, states the following:

> *Defensive operations defeat an enemy attack, buy time, economize forces, or develop conditions favorable for offensive operations. Defensive operations alone normally cannot achieve a decision. Their purpose is to create conditions for a counteroffensive that allows Army forces to regain the initiative. [5]*

These definitions of defense do not sound like our current approach to NETOPS and CND. Clausewitz might say we have a shield mentality about cyber defense. The O&M activities that we routinely refer to as 'network defense' are passive and do not try to gain or maintain the initiative. An active defense—one that employs limited offensive action and counterattacks to deny the adversary—will be required to have a genuinely defensive capability in cyberspace.

A Force Protection Model

So if NETOPS isn't CND, then what is it? Joint Publication (JP) 6-0, Joint Communications System, is the DoD's capstone document for communications and network support to joint operations. Chapter IV discusses NETOPS in depth, stating:

- **The effectiveness of NETOPS** is measured in terms of availability and reliability of network enabled services, across all areas of interest, in adherence to agreed-upon service.
- **The purpose of NETOPS** is assured system and network availability, assured information protection, and assured information delivery. [6]

The overarching theme in these statements is the ability for users (customers) to accomplish their missions, which leads us to the concept of 'mission assurance.' Mission assurance includes a number of activities and measures taken to ensure the availability of required capabilities and supporting infrastructures to support military operations and accomplish assigned missions. This includes areas such as force protection, antiterrorism, information assurance, and continuity of operations. [7] The security portion of NETOPS then can be viewed as a form of force protection, where force protection is defined as follows:

> *Preventive measures taken to mitigate hostile actions against DoD personnel (to include family members), resources, facilities, and critical information. Force protection does not include actions to defeat the enemy or protect against accidents, weather, or disease. [8]*

This definition does not say anything about defense in terms of maneuver and fires, but it does highlight that everyone in the DoD has a role in 'mitigating hostile activities' that can certainly be extended to cyberspace. There are a several reasons we should look at force protection doctrine as it relates to the NETOPS/security problem. The first is that force protection activities and doctrine are well-defined, and force protection experts have developed a rigorous methodology to define the force protection process, as illustrated in Figure 2.

The following force protection core principles apply to cyberspace:

Figure 2 Force Protection Planning Process [9]

- Determine the threat *via* a tailored threat assessment
- Determine critical infrastructure *via* a criticality assessment
- Determine vulnerability *via* a vulnerability assessment
- Determine acceptable risk *via* a risk assessment
- Develop a comprehensive force protection plan
- Exercise the plan to determine limiting factors and gain process familiarity.

A second reason to look at force protection is that force protection is an inherent responsibility of command. Air Force Doctrine Document 2-4.1, *Force Protection,* clearly states, "Commanders at all levels must make force protection an imperative." [10] A fundamental premise within JP 6-0 is that many of the responsibilities for NETOPS activities remain within the purview of the communications community. With a force protection mindset, responsibility shifts to the person who is accountable for mission accomplishment—the commander. At all levels of warfare, the commander should have the best understanding of both the mission and the requirements to accomplish it. The unit commander is therefore integral to cyberspace force protection actions and is not merely a customer. This conceptual shift integrates cyberspace force protection at the lowest possible level, thereby making it a unit commander's responsibility—which is where General Chilton said it should be!

Finally, the concept of force protection brings with it responsibility to every member of the force. The gate guards may "let the right people come in and keep the wrong people out," but we must be on the lookout for those who have gotten past the perimeter fence and those insiders who engage in malicious acts. Using a force protection paradigm, information assurance would equate closely to the Air Force (AF) Office of Special Investigations (OSI) 'Eagle Eyes'

construct. The AF OSI Eagle Eyes website states:

> *The Eagle Eyes program is an Air Force anti-terrorism initiative that enlists the eyes and ears of Air Force members and citizens in the war on terror. Eagle Eyes teaches people about the typical activities terrorists engage in to plan their attacks. Armed with this information, anyone can recognize elements of potential terror planning when they see it. [12]*

Conclusions

Semantics matter. One of the fundamental purposes of joint doctrine is to provide a common language that describes how we organize, train, equip, and employ our military capabilities. Inadequate semantics creates confusion and degrades our warfighting capability. Our current language confuses the use, operations and maintenance, and the defense of the cyberspace domain, which makes roles and responsibilities unclear. Our recommendations to remedy this situation are as follows:

1. Redefine NETOPS as "actions taken to provision and maintain the cyberspace domain." This would capture the current concepts of operations and maintenance while removing the ambiguity caused by including defense within the NETOPS construct.
2. Leverage concepts such as 'mission assurance' and 'force protection' to help change the culture and engage all personnel—users, maintainers, and cyber operators. Everyone has a role in security and force protection, but we are not all cyber defenders. Force protection and mission assurance are focused inward on our mission.
3. Redefine our CND construct to be more consistent with our approach to the concept of 'defense' in the other domains of warfare, to include the concept of active defense. This would shift the concept from maintenance to operations, from inward to outward (to our adversaries). CND is about delivering warfighting effects (*e.g.,* denying, degrading, disrupting, and destroying the cyber capabilities of our adversaries).

Taken together, these concepts provide a framework to develop cyberspace capabilities and personnel to meet joint mission requirements and to more effectively engage in operations in cyberspace. ■

About the Authors

Dr. Robert F. Mills | is an Associate Professor of electrical engineering at the Air Force Institute of Technology (AFIT), Wright-Patterson AFB, OH. He teaches graduate courses and leads sponsored research in support of AFIT's cyber operations and warfare program. His research interests include network management and security, communications systems, cyber warfare, and systems engineering. He retired from active duty in the US Air Force after serving 21 years as a communications officer.

Major Michael B. "Bo" Birdwell | is a career intelligence officer. He is the Director of Operations at the Air Mobility Command Air Intelligence Squadron at Scott Air Force Base, IL. Major Birdwell is a graduate of the Air Force Academy (1996), the USAF Weapons School Intelligence Division (2001), and the AFIT's Cyber Warfare Intermediate Developmental Education Program (2009).

Major Kevin Keller Beeker | is now the J2 Targeting Chief for the Joint Functional Component Command for Network Warfare (JFCC-NW) at Ft Meade, MD. He is a senior A/OA-10 combat pilot, who also completed an exchange tour flying F/A-18s with the United States Navy. He is a 1996 graduate of the United States Air Force Academy, with a Bachelor of Science in computer science. He is also a 2009 graduate of AFIT's Cyber Warfare Intermediate Developmental Education Program.

References

1. General Kevin Chilton, Opening Remarks to the April, 2009, USSTRATCOM Cyberspace Symposium, *http://www.stratcom.mil/speeches/23*
2. General Norton A. Schwartz, Letter to All Airmen, dated 27 May, 2009.
3. DoD Dictionary of Military Terms, *http://www.dtic.mil/doctrine/dod_dictionary*
4. Taken from Peter G. Tsouras. Warriors Words: A Quotation Book. 1992. Arms and Armour Press, London. Page 128.
5. US Army Field Manual (FM) 3-0, Operations, 14 Jun 2001, p. 1-15, *http://www.dtic.mil/doctrine/jel/service_pubs/fm3_0a.pdf.*
6. oint Publication (JP) 6-0, Joint Communications System, 20 Mar, 2006, p IV-1, *http://www.dtic.mil/doctrine/new_pubs/jp6_0.pdf.*
7. DoD Directive 3020.40, Defense Critical Infrastructure Program, 19 Aug, 2005, p. 13, *http://www.dtic.mil/whs/directives/corres/pdf/302040p.pdf.*
8. DoD Dictionary of Military Terms.
9. DODI 2000.16, DoD Antiterrorism (AT) Standards, provides clear guidance on the tools necessary to define the threat, determine what is critical, determine what is vulnerable, determine acceptable risk, develop a plan, exercise the plan, and then start over. The AT Risk Management process is outlined in enclosure 3 (pages 13—22). Available at *http://www.dtic.mil/whs/directives/corres/pdf/200016p.pdf.*
10. Air Force Doctrine Document 2-4.1, 9 Nov 2004, p. 11.
11. *http://www.e-publishing.af.mil/shared/media/epubs/AFDD2-4.1.pdf.*
12. The USAF OSI Eagle Eyes website is *http://www.osi.andrews.af.mil/eagleeyes/index.asp.*

The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

LPS-Public: Secure Browsing and an Alternative to CAC Middleware

by Lt Col Ken Edge and Kevin Sweere

On January 15, 2010, the Air Force Portal started granting access only to those users who have a Common Access Card (CAC) or public key infrastructure certificate, blocking login *via* user/password. Other Department of Defense (DoD) sites require CACs for some activities and it is likely many other federal agencies will also soon require two-factor authentication for sensitive Web services.

The DoD's solution for users of Windows XP Pro and Vista (a Windows 7 solution is coming soon) is to download licensed ActivClient middleware from an internal website. Users must install smartcard drivers, the middleware, and DoD root certificates on their Windows Personal Computers (PC). But that leaves out those running Mac or Linux systems, those using another's computer (*e.g.,* friend's, corporate or public computer), those lacking administrator privileges, and those who just do not want to make the requisite changes to update their computers. Lightweight Portable Security, Public edition (LPS-Public) alleviates all these problems. And it's free from *http://spi.dod.mil/*.

LPS-Public offers other benefits; computers that are old, slow, infected, or crashed, or those that are missing a hard drive can now browse the Internet again. Because LPS-Public operates only in Randon Access Memory (RAM), users may visit risky, malware-infected sites with very little permanent risk. Likewise, user's private sessions and sensitive transactions occur within a leave-no-local-trace browsing environment.

LPS-Public provides a thin, secure, end-node for cloud computing. Created by the Software Protection Initiative at the Air Force Research Laboratory (AFRL), LPS-Public boots from a CD, runs only in RAM, installs nothing to the hard drive, and does not require administrative rights. LPS-Public provides a Firefox browser with plug-ins, CAC middleware, certificates, and a PDF viewer within a very thin Linux operating system. It's a great solution for users with Mac, Linux, or Windows 7 systems, or those using others' computers.

A derived and accredited version, LPS-Remote Access, offers teleworkers remote desktop virtualization of their company's or agency's network. This means far fewer government laptops. Now one only needs to carry a CAC-reader and a custom CD and then use almost any personal, public, or corporate computer to use a NIPRNet computer remotely.

The Software Protection Initiative (SPI) protects critical DoD intellectual property against nation-state class threats by taking an alternative approach to security based on 3 Tenets: 1) Focus on What's Critical, 2) Move it Out-of-Band, and 3) Detect, React, Adapt. SPI solves your toughest cyber-defense challenges. The AFRL's ATSPI Technology Office manages SPI for the DDR&E *via* the High Performance Computing and Modernization Program.

Download the free LPS-Public ISO image from *http://spi.dod.mil/lipose.htm*.

Those wishing to get more details or interview a subject matter expert please contact Josh Aycock, 88 ABW/PA, at Joshua.aycock@wpafb.af.mil or 937-522-3514. ■

About the Authors

Lt Col Kenneth Edge | graduated from the US Air Force Academy with a degree in electrical engineering. His previous assignments in the Air Force have included flying C-141 and C-21 airplanes. Lt Col Edge completed his Master's degree in electrical engineering at Wright State University, and then earned his PhD in computer security from the Air Force Institute of Technology. He serves at the AFRL as the Office of the Director, Defense Research and Engineering's SPI Program Manager.

Kevin Sweere | serves the SPI as an Advisory and Assistance Services contractor from the not-for-profit Riverside Research Institute. He holds an Master's degree in Mechanical Engineering from Michigan Technological University and an MBA from University of Cincinnati. He was a search and rescue dog trainer, snowplow researcher, Army Ranger, Armor Battalion S4, satellite operator, and designer/builder of two bleeding-edge intelligence production centers. He now teaches his Tiger Scout den land navigation and fire building.

FREE Products 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 online: *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 _____

_____ Email _____

_____ Fax _____

Please check one: ☐ USA ☐ USMC ☐ USN ☐ USAF ☐ DoD ☐ Industry ☐ Academia ☐ Government ☐ Other

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

_____

LIMITED DISTRIBUTION

IA Tools Reports ☐ **Firewalls** ☐ **Intrusion Detection** ☐ **Vulnerability Analysis** ☐ **Malware**

Critical Review and Technology Assessment (CR/TA) Reports
- ☐ Biometrics (soft copy only)
- ☐ Configuration Management (soft copy only)
- ☐ Defense in Depth (soft copy only)
- ☐ Data Mining (soft copy only)
- ☐ IA Metrics (soft copy only)
- ☐ Network Centric Warfare (soft copy only)
- ☐ Wireless Wide Area Network (WWAN) Security
- ☐ Exploring Biotechnology (soft copy only)
- ☐ Computer Forensics (soft copy only. DTIC user code MUST be supplied before these reports will be shipped)

State-of-the-Art Reports (SOARs)
- ☐ Measuring Cyber Security and Information Assurance
- ☐ The Insider Threat to Information Systems (soft copy only. DTIC user code MUST be supplied before these reports will be shipped)
- ☐ Software Security Assurance
- ☐ A Comprehensive Review of Common Needs and Capability Gaps
- ☐ IO/IA Visualization Technologies (soft copy only)
- ☐ Modeling & Simulation for IA (soft copy only)
- ☐ Malicious Code (soft copy only)
- ☐ Data Embedding for IA (soft copy only)

UNLIMITED DISTRIBUTION

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

Volumes 11	☐ No. 1	☐ No. 2	☐ No. 3	☐ No. 4
Volumes 12	☐ No. 1	☐ No. 2	☐ No. 3	☐ No. 4
Volumes 13	☐ No. 1			

SOFTCOPY DISTRIBUTION

The following are available by email distribution:

- ☐ IADigest
- ☐ IA/IO Scheduler
- ☐ Research Update
- ☐ Technical Inquiries Production Report (TIPR)

Fax completed form to IATAC at 703/984-0773

Calendar

May

DISA Customer Partnership Conference
3–7 May 2010
Nashville, TN
http://www.disa.mil/conferences/

New York Metro Information Security Forum
4–5 May 2010
New York, NY
http://www.ianetsec.com/forums/calendar.html

Joint Warfighting 2010
11–13 May 2010
Virginia Beach, VA
http://www.afcea.org/events/jwc/10/intro.asp

IEEE Symposium on Security and Privacy
16–19 May 2010
Oakland, CA
http://oakland31.cs.virginia.edu/index.html

June

Forum of Incident Response and Security Teams (FIRST) Annual Conference
13–18 June 2010
Miami, FL
http://conference.first.org/

Lone Star Information Security Forum
23–24 June 2010
Dallas, TX
http://www.ianetsec.com/forums/calendar.html

July

2010 Software Protection, IA and Anti-Tamper SBIR Workshop
20–22 July 2010
WPAFB, OH
http://www.spi.dod.mil/workshop.htm

Black Hat USA 2010
24–29 July 2010
Las Vegas, NV
http://www.blackhat.com/html/events.html

DEF CON 18
30 July–1 August 2010
Las Vegas, NV
https://www.defcon.org/

August

LandWarNet 2010
3–5 August 2010
Tampa, FL
http://events.jspargo.com/lwn10/Public/MainHall.aspx

Air Force Information Technology Conference (AFITC 2010)
30 August–1 September 2010
Montgomery, AL
http://www.mc2-afitc.com/

To change, add, or delete your mailing or email address (soft copy receipt), please contact us at the address below or call us at: 703/984-0775, fax us at: 703/984-0773, or send us a message at: iatac@dtic.mil

IATAC
Information Assurance Technology Analysis Center
13200 Woodland Park Road, Suite 6031
Herndon, VA 20171