11/2

Volume 11 Number 2 • **Summer 2008**

IAnewsletter

The Newsletter for Information Assurance Technology Professionals

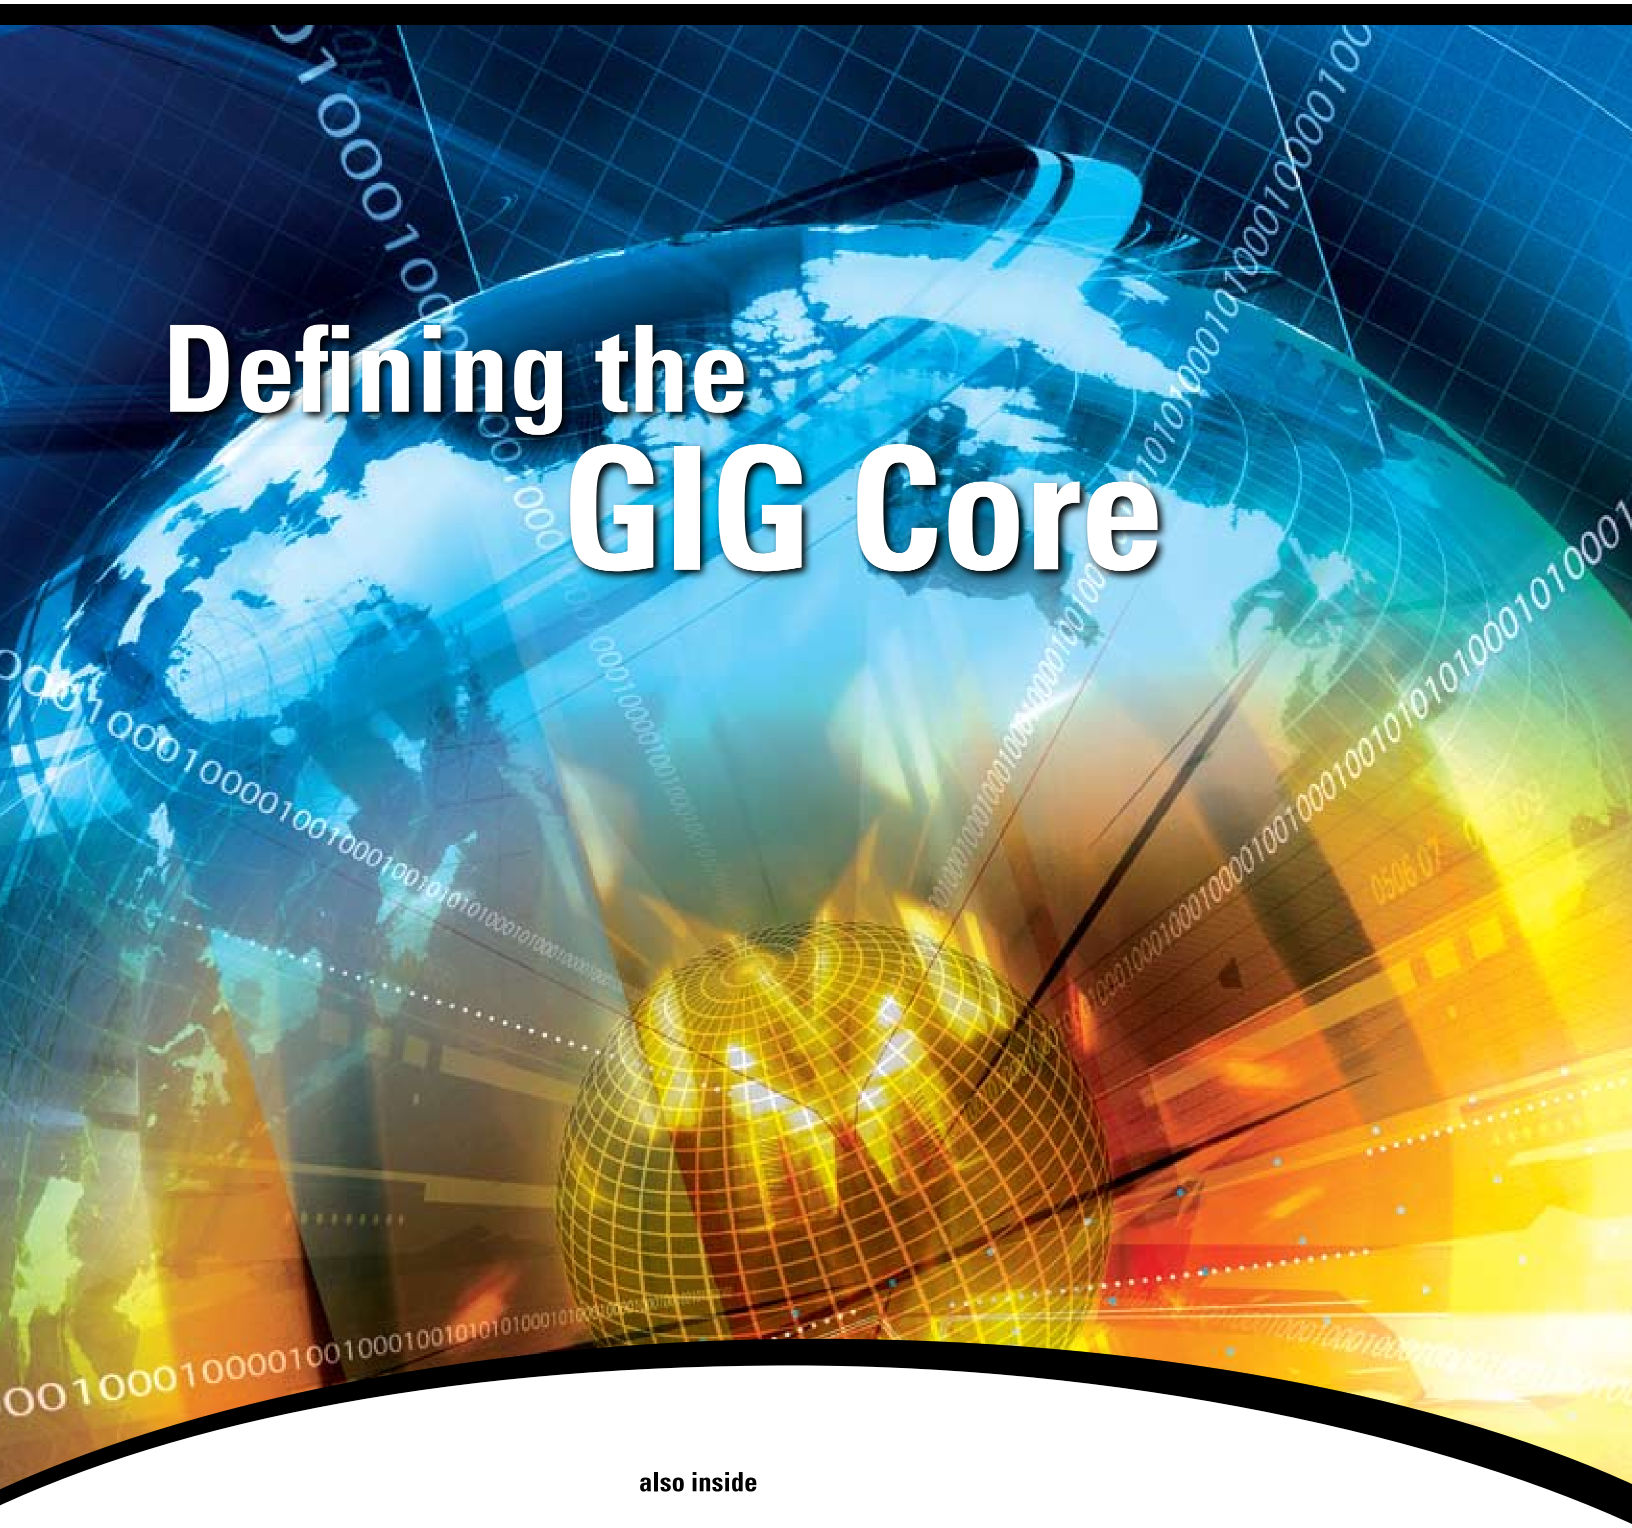

Defining the GIG Core

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 and/or 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

***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
Fall 2008 September 12th, 2008

Cover design: Dustin Hurt
Newsletter
design: Bryn Farrar
Donald Rowe

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

in every issue

IATAC Chat

Gene Tyler. IATAC Director

Often, organizations focus on trying to keep out unauthorized individuals, but they fail to notice those on the inside—individuals in positions of trust.

The Information Assurance Technology Analysis Center (IATAC) is pleased to announce the release of our much anticipated State-of-the-Art Report (SOAR), *The Insider Threat to Information Systems*. As I write this article, the SOAR is undergoing final revisions and approval; therefore, by the time this edition of the *IAnewsletter* is released, the SOAR *Insider Threat to Information Systems* should already have hit the streets.

This SOAR takes a close look at an often overlooked threat: the insider. Often, organizations focus on trying to keep out unauthorized individuals, but they fail to notice those on the inside—individuals in positions of trust. This report considers an "insider" as a person within an organization who is entrusted with privileges for accessing the organization's information, information systems, and/or facilities. This SOAR provides a comprehensive examination of the current state-of-the-art in addressing the insider threat, specifically as it pertains to information technology (IT) systems.

The report begins with an overview of how the insider threat is defined and viewed within government, industry, and academia. We review various policy, technical, and procedural approaches being applied across these communities to address the threat and then describe ongoing research that is meant to broaden our ability to limit or prevent an insider attack. Finally, the SOAR presents state-of-the-art best practices that government and industry are using to mitigate insider threats.

If you have never considered an insider attack as a possible threat, this report will make you rethink your position. I encourage each of you to reach out to IATAC for a copy of *The Insider Threat to Information Systems*.

IATAC also is pleased to announce the newly appointed Deputy Director for Research and Academic Integration and PhD candidate, Ms. Angela Orebaugh. Angela is an information security technologist, scientist, and author, with a broad spectrum of expertise in information assurance (IA). She synergizes her 15 years of hands-on experience within industry, academia, and government to advise clients on IA strategy, management, and technologies. Angela leads several security initiatives with the National Institute of Standards and Technology (NIST), including technical special publications (800 series), National Vulnerability Database (NVD), Security Content Automation Protocol (SCAP), and secure eVoting.

Angela is also an adjunct professor at George Mason University, where she conducts research and teaches intrusion detection and forensics. Her research includes peer-reviewed publications addressing intrusion detection and prevention, data mining, attacker profiling, user behavior analysis, and network forensics. Angela not only has authored several *IAnewsletter* articles but also is the author of the Syngress best seller's Nmap in the Enterprise, Wireshark and Ethereal Network Protocol Analyzer Toolkit, and Ethereal Packet Sniffing. In addition, she has co-authored the *Snort Cookbook*, *Intrusion Prevention and Active Response*, and *How to Cheat at Configuring Open Source Security Tools*. She is a frequent speaker at various security conferences and technology events, including the SANS Institute and the Institute for Applied Network Security.

Angela holds a masters in computer science and a bachelors in computer information systems from James Madison University. As mentioned above, she is currently completing her dissertation for her PhD at George Mason University, with a concentration in information security. IATAC is excited to have Angela as a welcomed edition to the IATAC team.

This *IAnewsletter* edition contains several intriguing articles. An IATAC subject matter expert wrote one such article, "Recent Developments in Cyberlaw." This article examines various technology and computer usage changes that have required current laws to be reviewed and new cyberlaws to be developed. Another interesting "must read" is, "Defining the GIG Core." This article reviews the intent of the Global Information Grid and what is still needed to achieve this vision. These are only a couple of many thought-provoking articles that you will find in this edition of the *IAnewsletter*. ■

Gene Tyler

FEATURE STORY

Defining the GIG Core

by Julie Tarr and Tony DeSimone

Abstract

This article defines numerous concepts associated with the GIG and discusses two architectural options for constructing the core of the GIG: striped core and black core. In all cases, we assume that traffic flows are protected in the core using Internet Protocol Security (IPSec) or similar protocols. A striped network simplifies the interconnection of core component by making traffic visible at the interconnection point, whereas decrypting and reencrypting to allow interconnection of core components complicate the end-to-end problem of IPSec gateway discovery, network routing, and quality of service. Decrypting at intermediate nodes also compromises the protection of traffic afforded by end-to-end IPSec encryption. We demonstrate that a black core offers greater flexibility in exploiting network connectivity than a striped core.

Introduction

The Global Information Grid (GIG) is a large, complex undertaking that is intended to integrate virtually all information systems, services, and applications in the US Department of Defense (DoD) into one seamless, reliable, and secure network. To achieve the GIG vision of ubiquitous and reliable communications, the GIG will need to support mobility, security, and survivability over a core network infrastructure that is built from network components that various services and organizations have procured and manage.

The network infrastructure is fundamental to the vision for future military operations and communications. [1,2] The capabilities needed for this future information grid (*e.g.,* mobility, security, and survivability) impose significant requirements on the network. This article discusses the architectural options in constructing the network, including the protection of traffic traversing the core, the interconnection of core components, and implications for quality of service and routing.

Needs and Challenges

The GIG will be diverse in not only the necessary technologies for supporting GIG capabilities but also the range of operational environments. This diversity is inherent to the missions. Future warfighters will require information at the edge, delivered over tactical wireless networks. The information includes local and regional communications, along with reachback over satellite to resources in data centers attached to high-speed terrestrial networks. Bandwidth of systems providing connectivity in this environment could span six orders of magnitude, from forward-deployed tactical systems at a few kilobits per second (kb/s) to fiber networks and attached resources at 10 Gigabits per second (Gb/s).

Further, the GIG is not a single program; rather, it is a construct for driving the development of multiple programs. The network infrastructure will be built from components of various services and organizations. Technologies will be diverse, policies will limit what can be communicated across network interfaces and operations that bring together various services, and organizations add complexity in the interconnection of the networks.

Network Core

The GIG network infrastructure will be an Internet Protocol (IP) based "network of networks" composed of network components that are controlled and managed by various organizations or services. Each core component will be administered separately. The network core components are also composed of various transport mediums (*e.g.,* fiber, wireless, and satellite links). The performance characteristics of each network core component may vary widely. The network infrastructure also requires interconnection to the Internet.

Because the network infrastructure is IP based, we assume that IPSec [3,4,5] devices will be used to protect information across the core of the network. IPSec establishes security associations (SA) between each pair of communicating entities. When IPSec is used as a gateway, each pair of IPSec devices will establish SAs for all

communication between their respective user networks. User networks are referred to as plaintext (PT) or **red networks**. A **black network** refers to a network that is transporting *only* IPSec-encrypted traffic or **black traffic**. Black networks are also referred to as ciphertext (CT) networks.

Two options should be considered regarding the encryption of traffic as it traverses the core. In the first option, the core is black as defined above. In a black core, unclassified traffic destined to the Internet will be encrypted as it traverses the core; thus, encryption/decryption gateways will be required at the interconnection points between the GIG core and the Internet. In the second option, only classified traffic will be IPSec encrypted. Unclassified traffic may or may not be encrypted. We use the term **unclassified network** for networks transporting any unencrypted unclassified traffic.

Connections Between Core Components

Two approaches are possible for the interconnection of core network components. In the first approach, all data remains IPSec encrypted as it is transported across the core. For simplicity, we assume that all traffic is IPSec encrypted, and we use the term **black core** to describe this black-to-black interconnection of core network components. In the second case, all data is decrypted and re-encrypted when passing from one core component to another. In this case, the core is **striped** and the encryption/decryption nodes are termed **red gateways**. Figure 1 presents a simplified view of a striped core. The network design implications of a black or striped core are discussed below.

Black Core

As discussed above, in a black core all data remains IPSec encrypted as it is transported across the core. Figure 2 illustrates a multiple component network with a black core. All core component networks are directly connected thus there is interworking of routing and quality of service across the core.

One advantage of a black core is the maintenance of end-to-end [6] security services across the core. The confidentiality, authentication, and integrity provided by the IPSec encryption applied by the owner of the data is maintained across the core, simplifying the trust relationship required with intermediate networks that process and store this information.

The primary disadvantage of a black core is in computer network defense (CND)—in particular, the difficulty in applying perimeter protection at the edge of each administrative domain when all traffic is encrypted. Current perimeter protection mechanisms (*e.g.,* firewalls, intrusion detection systems, and virus scanning) require that all user data be decrypted. In addition, tactical networks may wish to filter traffic to reduce the load on disadvantaged links and/or networks. Filtering of encrypted traffic based on anything other than source/destination is difficult when traffic is encrypted.

Another possible disadvantage of a black core is the scale of the resulting network. The routing architecture will need to be designed to address the overall scale of the core. Scalability may be an issue with discovery mechanisms for IPSec encryption devices. [7,8] Coordination also will be needed in the overall design of the network architecture. Each core component is a part of the black core; the network architecture, including routing and quality of service (QoS), will need to be implemented consistently across the entire network.

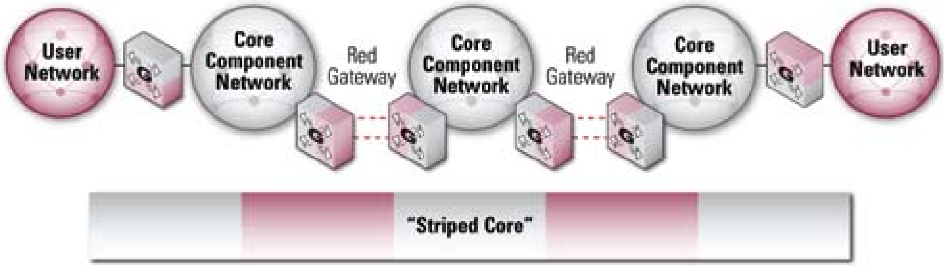

Figure 1 Striped Core

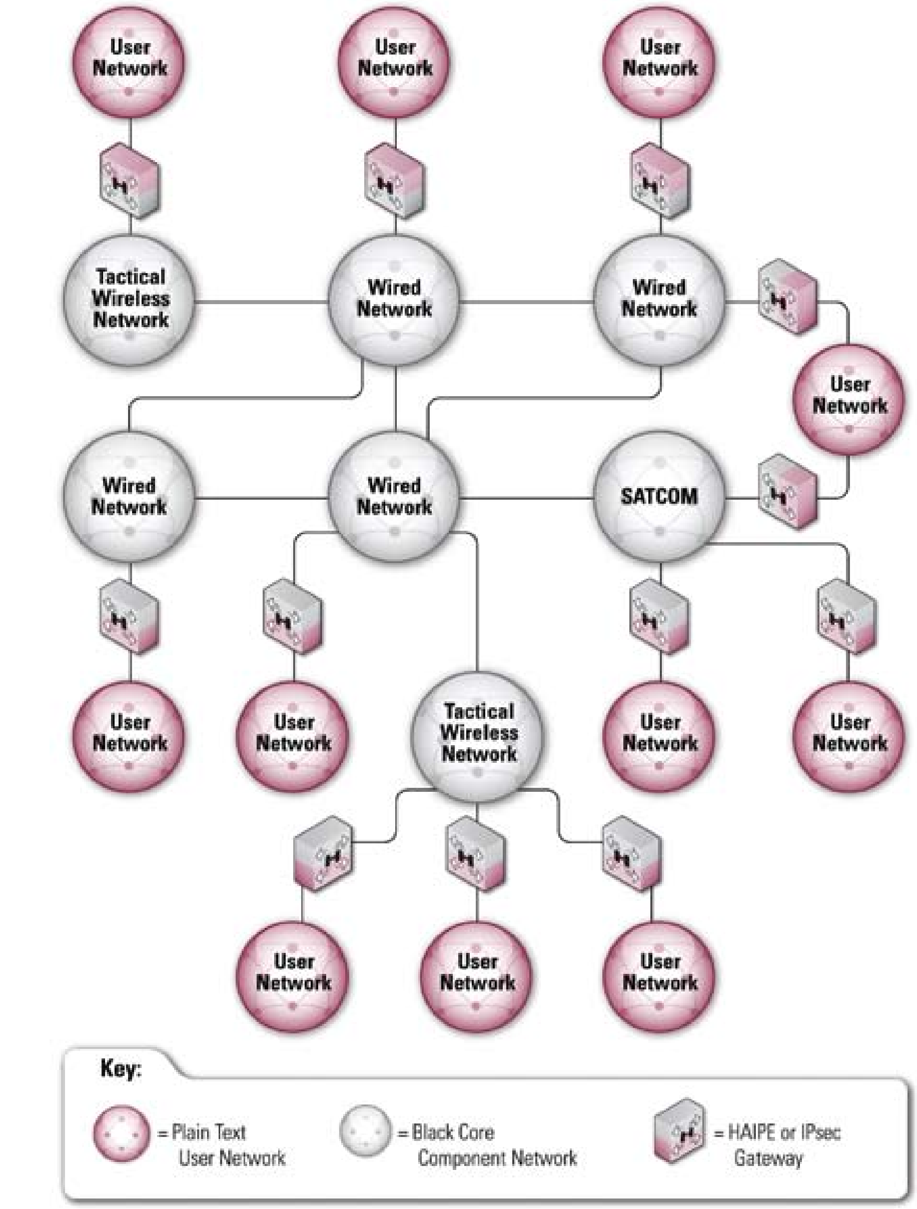

Figure 2 Multiple Component Network With Black Core

Striped Core

In a striped core, all data is decrypted and reencrypted at red gateways when passing from one core component to another. Figure 3illustrates a striped core.
For simplicity, we have shown only user networks for a single security level.
In reality, the network will be required to support communication for many security levels. Multiple gateways, one for each security level of data being transported across the network, will be required at each interconnection of core component networks. Again for simplicity, we will discuss user networks and red gateways operating at a single security level.

Red gateways impose restrictions on the operation of the overall network. IPSec devices envisioned for deployment in the GIG cannot pass routing information; thus, red gateways isolate each core component networks from the remainder of the network. The red gateways also break the end-to-end security services (*e.g.,* confidentiality, authentication, and integrity) provided by the IPSec encryption under the control of the owner of the data. All intermediate networks must be trusted to process and store all data being transported through their network, and a red gateway must be established for every security level carried by that intermediate network, which multiplies the cost and complexity of the interconnection. The dynamic formation of communities of interest (COI) will be hindered by the complexity of establishing red gateways, across all intermediate networks, at the new security level of the COI.

A striped core allows existing CND approaches to be maintained through the application of perimeter protections at administrative boundaries based on all information in the traffic flow. Stripes also mitigate some scalability issues: IPSec gateway discovery is now limited to the IPSec gateways in each administrative domain, limiting scalability requirements on any protocol developed to support IPSec gateway discovery. That benefit carries a cost: as discussed in the next section, the IPSec gateway discovery process cannot identify all IPSec gateway options that are available end to end because information from other core components is unavailable.

Routing

All aspects of performance, including routing, will be affected by the architectural decisions regarding interconnection of core component networks, described above as black core and striped core. The interconnection must be designed carefully to enable any-to-any communications to support the requirements of the GIG.

Routing in the GIG is broken into two parts: routing between endpoints and an IPSec gateway (*i.e.,* red-side routing), and routing across the core, between IPSec gateways. The IPSec gateway discovery mechanism addresses the mapping between red-side routing and routing across core. This section discusses the implication of the type of core, black versus striped, on the IPSec gateway discovery process. (The previous section discussed issues regarding routing across the core for black and striped cores.)

In the black core, shown in Figure 2, only one pair of IPSec gateways is used for a traffic flow. The ***IPSec gateway discovery problem*** is solved once across the black core, and routing across the core is completely determined by the operation of the routing protocol. In a striped core, shown in Figure 3, a pair of IPSec gateways protects the traffic across each core component; therefore, any traffic flow will be decrypted and reencrypted at each stripe. In cases of interest for the GIG, routing information does not propagate across the IPSec gateway, and the IPSec gateway discovery problem is solved serially at each stripe. The destination information is mapped to an IPSec gateway at each stripe, so providing a path as good as the path across the black core requires that the gateway discovery protocol have knowledge of routing sufficient to make the right choices of IPSec gateways.

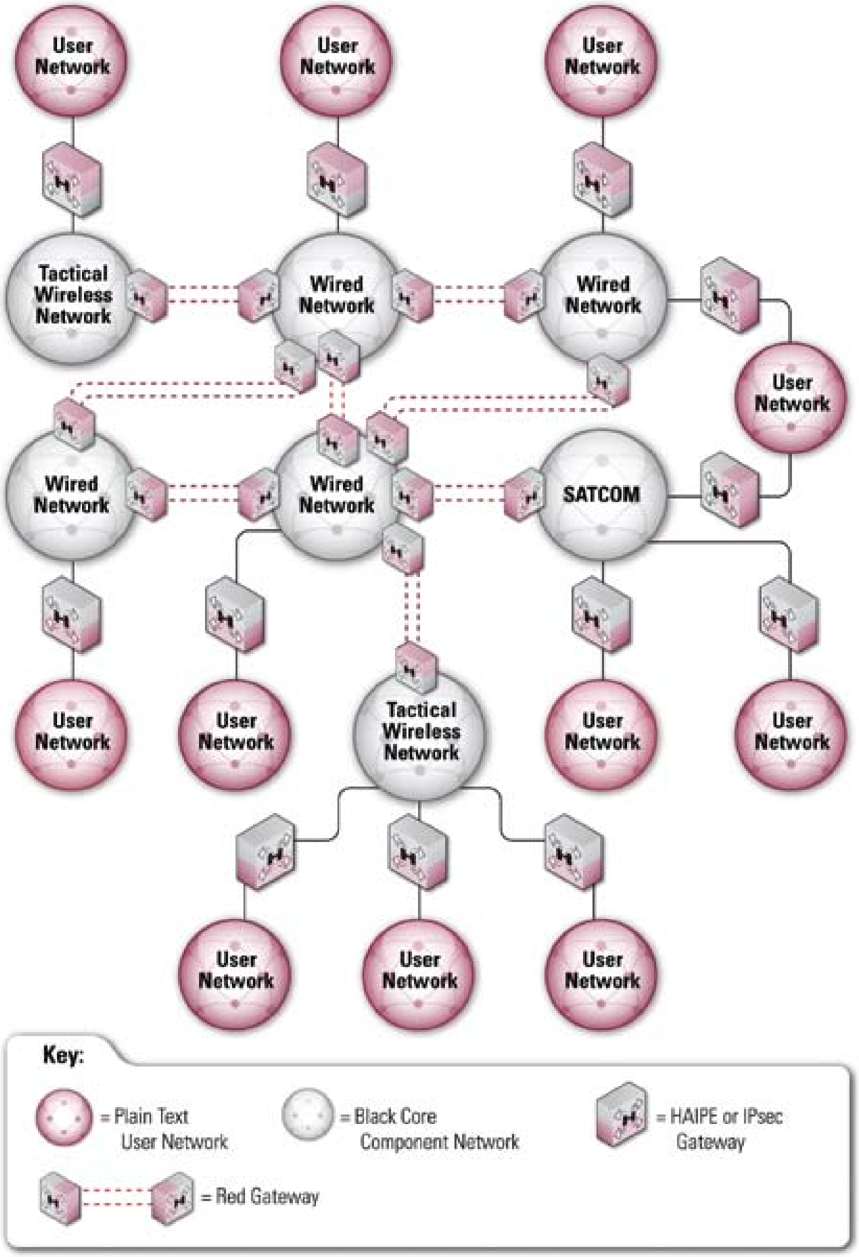

Figure 3 Multiple Component Network With Striped Core

Because IPSec devices are used to protect information across the network, the network must have a robust means of identifying the IPSec gateway associated with a destination address. The solution to this IPSec gateway discovery problem will determine how effectively network resources can be used.

Communications among users requires IPSec sessions to be established across each core stripe. Each IPSec gateway maintains a correspondence table that identifies a set of PT addresses reachable *via* a CT address. We will assume that the IPSec gateway can create a correspondence table for hosts on the locally attached PT network. For example, this effort could be achieved by participating in routing on the PT side. The correspondence table must be exchanged with other IPSec gateways attached to the same CT network. Numerous approaches have been used to propagate the correspondence tables, the most common of which is a full mesh of connections across the CT network; however, more scalable solutions are possible. For PT addresses on other CT networks, either a static IPSec gateway configuration or a default IPSec gateway CT address is needed.

Figure 4 shows an interconnection of core components *via* stripes, with IPSec gateways. As shown in the figure, when endpoint A sends traffic to endpoint B, IPSec gateway H1 will have a table that identifies the corresponding IPSec gateway to communicate with B: H3 in this case. A single SA [9] between H1 and H3 is sufficient to enable A-B communications. When endpoint A sends traffic to endpoint C, none of the IPSec gateways on the CT network directly connect to the PT network for endpoint C; therefore, H1 establishes an SA to the default IPSec gateway: H2 in this example. H2 decrypts and sends the traffic to H4, its matching IPSec gateway across the red stripe, which carries the correspondence table identifying H5 as the correct IPSec gateway for endpoint C. A and C can now communicate across the concatenated security associations in the CT networks and the PT stripes.

In Figure 4, routing across the CT networks is always based on the CT address of the IPSec gateway, which is determined by the IPSec gateway discovery process. Routing across the PT networks, however, is based on the PT address of the destination. Each PT component of the striped core network, or red gateway, must manage its routing architecture to ensure that a packet destined to any PT address outside that component network will be routed to an appropriate IPSec gateway.

A few approaches are possible: the IPSec gateways could redistribute routes into the PT segment for other PT addresses, static routes for some set of PT prefixes could be maintained, or default routes could be configured and managed in the PT segment to route packets to the appropriate IPSec gateway. It is likely that a combination of these approaches would be used, depending on the complexity of the PT segment. Regardless of the approach, the complexity and operational overhead associated with maintaining a PT routing architecture increases if the architecture needs to support routing through a striped core.

Further, the approach for striped cores limits how connectivity can be used. In the example of Figure 4, if H4 is unable to communicate with H5, no communication is possible between A and C. The concatenated SAs H1-H3 and H6-H7 would not be found by the techniques described above.

If the stripes are eliminated and the interconnection of component networks is made through black interfaces, any communications can be established thorough a single SA. The exchange of correspondence tables is now among a larger set of IPSec gateways, increasing the importance of a scalable approach for IPSec gateway discovery. Once the correct IPSec gateway is identified, however, the full capabilities of routing are used to identify a path to the correct IPSec gateway, including routing around failures if parts of the network are compromised.

Quality of Service

QoS mechanisms will operate in each component network that comprise the core, under the control and management of the respective administrative authority. To deliver services end-to-end, the QoS mechanisms will need to operate across the interfaces between the component networks and across the IPSec devices.

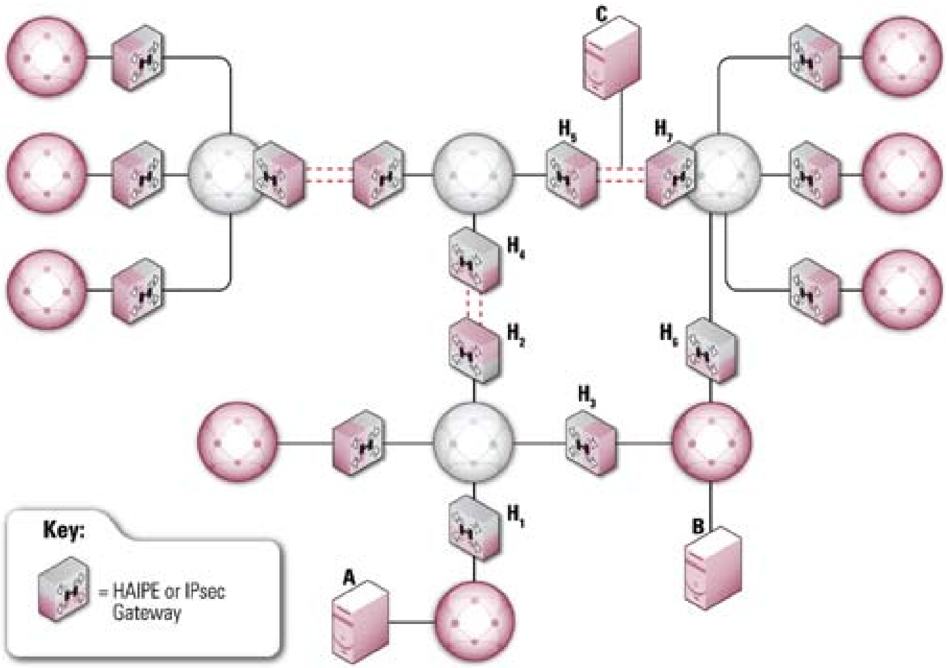

Figure 4 IPSec Gateway Discovery Across Red Stripes

The component networks will be administered separately; thus, coordination of QoS mechanisms across the component networks will be required to provide end-to-end QoS consistent with the user's performance requirements. For example, if packet markings are used to indicate the user's performance requirements, to meet those requirements end-to-end, the separate administrations will need to agree on the interpretation of markings in terms of performance and will need to implement mechanisms in their networks that support the user's needs. If signaling is used to support QoS, the component networks must be able to consistently interpret signaling and use the signaling to implement control mechanisms that support the user's QoS requirements.

Because IPSec devices are meant to protect information across the core of the network, the QoS information that crosses the IPSec device will likely be limited. The IPSec gateway therefore represents an impediment to providing end-to-end QoS, including more IPSec gateways as in a striped core, and complicates the QoS problem.

Conclusion

The network infrastructure to support the GIG will be based on an interconnection of components and will need to protect traffic as it traverses the network core. We have described two architectural approaches for such a network that provide different tradeoffs of risk, performance, and management complexity.

To support the GIG's mission, the network infrastructure must be designed as a whole. The dominant issue is the approach to the core infrastructure: black versus striped. Careful consideration needs to be taken in implementing stripes in the network. Although providing a means to control traffic across administrative domains, stripes complicate and limit other capabilities (*e.g.,* IPSec gateway discovery, routing, and QoS), creating a brittle core. The CND controls implemented at the stripes also become less effective as the protection model changes. Even when IPSec is used at the network layer, other encryption techniques are used end to end. Decrypted network-layer traffic may still be encrypted at the session or application layer (*e.g.,* Secure Sockets Layer [SSL] or data object encryption), thereby invalidating all filtering techniques that look at higher layers.

Although interconnecting components in the black raised new challenges regarding how we think about service delivery at network boundaries and about CND, the benefits in agility, simplicity, and economy point to the value of minimizing stripes and driving to an architecture based on a black core. ■

References

1. Alberts, David S., Richard E. Hayes, *Power to the Edge*, International Standard Book Number (ISBN) 1-893723-13-5, UB212.A43, 2003.
2. Alberts, David S., John J. Gartska, Frederick P. Stein, *Network Centric Warfare*, ISBN 1-57906-019-6, U21.2.A413, 1999.
3. Kent, S. and R. Atkinson, *Security Architecture for the Internet Protocol*, Request for Comment (RFC) 2401, November 1998.
4. Kent, S. and R. Atkinson, *IP Authentication Header*, RFC 2402, November 1998.
5. Kent, S. and R. Atkinson, *IP Encapsulating Security Protocol (ESP)*, RFC 2406, November 1998.
6. In most cases, the security services provided by IPSEC encryption are not truly end-to-end, rather IPSec-gateway-to-IPSec-gateway. For readability, we will use the term "end-to-end."
7. Small, S., Antonio DeSimone, Bharat Doshi, Fabian Monrose, Andreas Terzis, *Large-Scale Dynamic Virtual Private Networks for the Global Information Grid*, MILCOM 2005.
8. Nakamoto, Glen, Lisa Higgins, Justin Richer, *Scalable HAIPE Discovery Using a DNS-like Reference Model*, MILCOM 2005.
9. To be strictly correct: a pair of unidirectional SAs. We will use the shorthand for readability.

About the Author

Julie Tarr | is a Senior Enterprise IA Systems Engineer in the Applied Information Sciences Department at the Johns Hopkins University Applied Physics Laboratory. Previously, Ms. Tarr was a Senior Systems Engineer in the GIG Enterprise-Wide Systems Engineering Office in OSD(NII)/DoD CIO. Ms. Tarr led the development of the GIG Technical Direction, the NCID, and was a principal member of the team developing the GIG Technical Foundation. Ms. Tarr led the GIG Routing Working Group defining the routing interoperability architecture for the GIG. Prior to her assignment at NII, Ms. Tarr was the Head of the Secure Network Section at the Naval Research Lab which performs research, development, and support to the Navy in the areas of network security, intrusion detection, cross domain solutions, and networking requirements for encryption devices. Ms. Tarr was a member of the NATO Task Group on NEC Security and the TTCP Technical Panel on IA. Ms. Tarr received her BS in Electrical Engineering and Mathematics from the University of Maryland, College Park, in 1985. As a Naval Research Laboratory Fellow, Ms. Tarr received her MS in Electrical Engineering in Communication and Control Theory from the University of Maryland, College Park. She may be reached at *julie.tarr@jhuapl.edu*.

Tony DeSimone | is the Deputy to the GIG Senior Systems Engineer in the office of the Assistant Secretary of Defense for Networks and Information Integration/DoD Chief Information Officer, where is working to implement the vision of the GIG. Since 2004, he has been the Senior Scientist in at the Applied Physics Laboratory, Johns Hopkins University, investigating network architectures, security, routing and network management in support of GIG programs. Prior to joining APL, he led the System Engineering and Integration team in Lucent Technologies Government Solutions. He has held senior management positions in Lucent's Optical Networking Unit and Lucent Digital Video. At AT&T WorldNet, he was a product manager responsible for delivering Internet Telephony services. At Bell Laboratories, he did work in network design, performance analysis, network and computer security, wireless networking and digital video. He holds ten patents in data networking, Web caching and other Internet applications and has authored numerous technical publications. He earned a PhD and ScM from Brown University and a BS from Rensselaer Polytechnic Institute, Troy, NY (1981), all in Physics. He may be reached at *antonio.desimone@jhuapl.edu*

Letter to the Editor

Q ***I am aware that the DoD has created a new US Command, but do not know anything about it, could you help?***

A On 6 February 2007, it was announced that the US Department of Defense was creating a new US Africa Command headquarters: AFRICOM. Before AFRICOM, the US military interaction throughout Africa had been assumed by the US European Command (EUCOM), US Central Command (CENTCOM), and US Pacific Command (PACOM). However, to support Africa with its continued desire to build democratic institutions and establish good governance across the continent, one consolidated command needed to be established.

General William "Kip" Ward, Commander, US Africa Command, states, "AFRICOM seeks to build partnerships to enable the work of Africans in providing for their own security. Our intent is to build mutual trust, respect, and confidence with our partners on the continent and our international friends." Currently AFRICOM is operating under EUCOM, but it is anticipated to be a fully operational command by October 2008.

For additional information, please visit the United States Africa Command website at *www.africom.mil*. ■

Tomorrow Night

by Maj Paul Williams, Dr. Steve Rogers, Dr. Robert Mills, Dr. Richard Raines, and Timothy Lacey

It is late night in the United States, with darkness stretched evenly between dusk and dawn. Most of the garrison US Air Force (USAF) slumbers, its airmen sleeping quietly. So does much of its information infrastructure. Very akin to the dreams of the airmen who make up the USAF, the networks that serve them also dream. Information about the previous day's operations flow throughout the networks awareness nodes, enhancing the network's sense of self, of normalcy, and of the ebbs and flows of data through its links. There is tension in the dreams of the network and its users—the United States and an adversary have been slipping closer to conflict. The news has been full of stories and speculation about whether or not the differences will have a peaceful resolution, and the daily activity of the USAF has been ramping up to support a possible conflict. The increased planning activity of its users represents change to the network, which it absorbs into its world model through dreamlike activity. This assimilation of recent experiences and the knowledge and past experiences embodied within the world model is a crucial enabler of its ability to anticipate future data flow requirements and threats. Still, continued peace seems almost certain, so the daytime USAF sleeps.

It may or may not be night where these adversaries live, where they plan their operations, or where they contemplate how to advance their causes. Their best interests, they believe, require taking a chance. They choose to act when the United States sleeps, counting on surprise to slow the giant's response enough that they will be able to achieve their objectives before the might of the United States' military crashes down on them. Knowing the US military's dependence on its information infrastructure, the adversaries had invested heavily in network warfare capabilities. Confident in their ability to disrupt our ability to communicate and operate, they are certain they can inflict strategic paralysis upon us. Long laid plans are activated, carefully practiced activities are executed, and at 0200 US Eastern Time, stealthily implanted capabilities awaken and strike from within even as waves of attack activity surges against the USAF's network perimeter defenses.

The USAF's network was dreaming, but this does not mean its eyes and ears were disconnected—in this regard, its slumber differs from that of its users. At 0200, a sudden surge of activity courses both outward and inward through the links separating it from the outside world. As the surges occur, sensors capture their essence and forward it to the self-awareness engines to wash over the network's model of itself. The behavior does not match anything the network expects—it does not feel familiar, it feels alarming! The sensor flows are also passed through known and suspected threat models, and some patterns match activity detected during the adversaries' training activities. Within a fraction of a second, automated responses are invoked, certain types of incoming and outgoing network traffic are severely curtailed, snapshots of the critical infrastructure state are taken, and alarms are set off in the USAF's watch centers to alert the cyber warfare forces of a possible attack.

By 0201, the network recognizes that it is in severe trouble and attempts to activate even more warnings to its human operators, but in most cases the adversaries' deeply embedded attack capabilities have severed connections among the far-flung parts of the USAF's enterprise. Each base is on its own, and its information infrastructure is under constant attack from within and without. The same is true across all the Department of Defense (DoD), government agencies, and the US critical infrastructure. We had been hit! Hard!

Fortunately, the now disconnected network nodes contain high-resolution models of the remainder of the network and can predict what is happening within those domains it cannot currently contact.

0205 Eastern Time. The adversary's embedded sensors report back *via* satellite that the initial stages of the attack appear to be successful. Knowing the reliance the US military has in its information infrastructure, and confident that the United States is now paralyzed, the senior decisionmakers authorize the kinetic phase of operations.

The network is fully alert now, and significantly enhances its introspective activities, scanning deeply into itself for anomalous activity. Awareness of the type, scope, and intent of the attack develops as the network begins to construct models of the malicious activity. These models lead to predictions about what is happening; the network activates unused sensors to validate its understanding about what is happening, and in some cases creates new sensors and deploys them on the fly.

The network is able to monitor and model itself, as well as activity within itself, across multiple axes—from network traffic flows to the behavior of low-level activity in the hardware and software of its nodes, and from users and their patterns of activity to organizations and the missions they support. Just as with humans, these predictions are in fact the ”reality” on which the network bases its actions—an approach that not only allows for quicker response, but also requires constant feedback, refinement, and adaptation.

As these models develop, they are presented to the cyber warfare officers and non-commissioned officers responsible for defending our information infrastructure. This information augments their understanding of the underlying attack behavior, dramatically increasing their situational awareness. They augment the automated responses and inward directed probing with their own questions and directed activities. Within several minutes, they are peeling back the layers of the attack, identifying which components are simply noise masking more directed activity. The patterns emerging lead to an understanding of the attacker's intent and thereby to the development and employment of effective countermeasures.

By 0230, adversary forces arrayed next to a national border, notionally for an exercise, surge into action. Reports from the network attack units continue to reflect success.

By now, the USAF's network warfare forces are beginning to reconnect those network segments disrupted by the attack. As the network awareness nodes merge into the larger network self, an even better understanding of the attack patterns emerges. Nodes with malicious code embedded in software and firmware are identified and automatically cleansed.

In places where attack capabilities are operating out of hardware, the corrupted nodes are isolated and captured for further study. The attacker's intent, methodologies, and report-back capabilities are now well enough understood that it is possible to feed deceptive information back to the attacker. Senior US officials order that this happens while traditional intelligence capabilities use the captured intent from the network warfare attack to focus their efforts on the kinetic battlefield.

By 0500, aircraft have launched and formed into attack groups; tanks have started up and are headed for the border. In both cases, US and Allied forces meet them in the air and on the ground.

All available intelligence begins reporting that the US military's operations are unimpeded. Stunned, the adversary's leadership orders a full retreat, claiming loudly that the attempted incursion was merely a part of the exercise. Furious, the senior leadership demands to know what had happened. Why had the crippling of our information infrastructure not prevented our response?

Using a combination of intrinsic network warfare capabilities and trained deception personnel, the robust understanding of the attacker's abilities and intent give us options: we can let them think they are still successfully attacking us, or we can cut them off at the knees. By letting them continue their offense, we map out their injection and report-back means and enable more precise offensive actions later.

The network warfare forces had correctly determined that their attack against the USAF's network had been successful and devastating. They attempt to reach back into the United States' networks using the same techniques used earlier, and continue to see what they expect from a successful attack. They do not understand why we are still able to function or how their kinetic forces were contained and defeated so easily.

Back in the United States, the USAF is fully awake and analyzing the intelligence gleaned from the night's attack.

Highly trained network defense warriors stationed throughout the USAF's network infrastructure had repulsed the network attack using two new capabilities: a self-aware, self-healing network infrastructure coupled with a skilled response. The network had learned the form of the attack within seconds, partly by observation and partly by routing offensive activities into sandboxes and interacting with them, exposing models of itself to those attacks, and learning the internal and the attacker's responses through many parallel trials. In addition to taking certain automatic defensive actions, the network provided that information to operational defensive forces. These forces used this information to enhance their situational awareness and were able to rapidly counter the attacks and restore connectivity throughout the infrastructure in time to detect and respond to the adversarial force movements.

Back to Today

The story above is science fiction about technology we may be using in the future. The story also describes how we will conduct defensive network warfare in the future. Although the threat described above is realistic, the portrayed ability to deal with such an attack is not yet available. The self-aware networks depicted are under exploration by USAF scientists, and the network warfare forces implied by the story are currently in the formative stages. Like much good science fiction, however, networks such as those may soon be reality. The USAF is well along the path of developing a new career force that includes personnel capable of performing as described. This article briefly describes both.

Self-Aware, Self-Healing Networks

The Air Force Institute of Technology (AFIT), in conjunction with the Air Force Research Laboratory (AFRL), is performing groundbreaking basic research into a new way to solve some very difficult problems in pattern recognition. Essentially, our current machine learning and artificial intelligence capabilities fall far short of our requirements across a wide variety of warfighting domains. Our need to automate the kill chain, Anticipate, Interact, Find, Identify, Fix, Target, Track, Engage, Assess, Anything, Anytime, Anywhere (AIFIFT2EA4)—particularly in cyberspace where the relevant targets operate at wire speeds and in ways not directly visible to humans—requires that we explore new ways of solving these problems. [1]

AFIT and AFRL are tackling this problem by exploring a new way of computing that mirrors in some ways how humans, animals, and plants recognize those things that are important to them, reason about them, and act based on the results of that recognition and reasoning process.

For IANewsletter readers, this should sound familiar; we published an initial version of this work in Vol. 10, Number 3. [2]

This article illustrates in a "Day in the Life" sense how those capabilities might be employed in conjunction with other developing technologies, and who might employ them. Because we are working in the context of the QUalia Exploitation of Sensor Technologies (QUEST) project, the current project is called Network QUEST (N-QUEST).

We are working toward a computing capability that is concept-based versus symbolically-based. To explain the difference, consider the task of recognizing the letters in the Completely Automated Public Turing Test for Telling Computers and Humans Apart (CAPTCHA) puzzle [3] in Figure 1.

Today's computing tools cannot reliably solve CAPTCHAs—but humans can easily solve them. To tie this concept to the network domain, consider traditional signature-based detection of malicious code, anti-virus systems. An anti-virus signature is created based on forensic analysis of a captured virus. To keep the false positive rate down—that is, the alarming by the signature on non-malicious code—the signatures are fairly specific. This means that when a new variant of an existing virus is created, the signatures for the parent virus often cannot detect the children. Another familiar analog is network-based intrusion detection that uses signatures of known-bad entities traversing a network to detect some malicious activities. As in the virus domain, new network-based attacks are often simply modified versions of older, known attacks. As in the virus domain, catching them all remains an unsolved problem.

Both of these are similar to the CAPTCHAs above, but the letters and numbers are twisted versions of the original symbols.

What is it about us that enables us to solve CAPTCHAs? Humans are able to easily recognize the numbers and letters in the image above because we have an internal representation of written symbols in our language that is based on the concept of that number or letter, not some mathematical description of the relationship between the dark pixels in the image. In other words, we have an abstract understanding of the symbols that enables us to easily recognize them, even when they are distorted as in the CAPTCHA.

We are associating this representation with the concept of qualia. We use this term in QUEST to distinguish the difference between the stimulus and the internal representation that nature uses for exploitation. It is the red you "see" or the pain you "feel" as distinguished from the photons impinging on your retina or the nerve pulses transmitting to your brain as

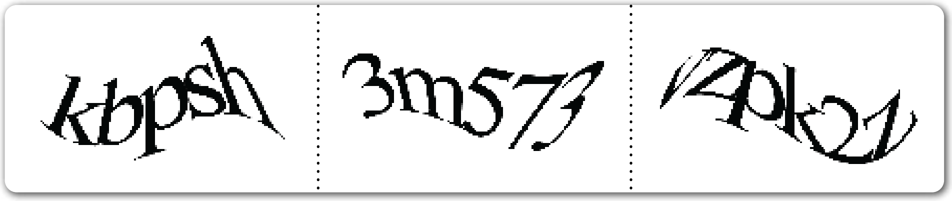

Figure 1 The CAPTCHA Puzzle illustrates a recognition task that is simple for a human but difficult to solve reliably using machine learning or artificial intelligence.

Tenet	Description
How will we know if N-Quest has Qualia?	A computer-based solution has formed a Qualia-like representation if the solution satisfies the set of tenets described in. [1] Our goal is to obtain an engineering advantage by using these Qualia-based tenets as a guide. Computer-based qualia representations are not assumed to be what animals experience.
Self	The concept of self involves being able to distinguish in the world model what is under your control and what is external. This computation will arise from interaction with the environment. N-QUEST will construct a world model based on its current environment and then place "itself" into that model.
Internally Generated	Qualia are evoked as a result of stimulation. That stimulation can be the result of sensing or can be internally generated, for example, by thinking or dreaming. Dreams are generated from qualia components. Dreams are our means to modulate our qualia representation based on recent experiences. We anticipate that N-QUEST will use a similar mechanism to assimilate sensed activity with past experience and knowledge.
Processes that Act on Qualia	A set of processes exist that manipulate the internal qualia representation. These processes generate efficient representations such as the formation of hierarchies to generate compound qualia. The processes of thinking and reasoning involve manipulating the links of which the qualia representation is composed. Processes manipulating qualia associated with confirmed predictions can be different from those for unconfirmed predictions.
Evolving Qualia	The qualia-based representation facilitates anticipating, detecting, distinguishing, and characterizing entities. That representation can change and be manipulated. The manipulation can be based on context and experience to identify the most suitable representation for an entity. Self entities in N-QUEST include network components, users, organizations, missions, and learned relationships between entities. Non-self entities include malware, network intruders, automated and human directed and illicit activity by trusted users. We anticipate tens or hundreds of thousands of such entities for even base-sized enclaves.
Qualia Theory of Relativity	Qualia-based representations build a world model that is completely relative. Each individual quale can be characterized only relative to other qualia (thus link-based representations are used for qualia).
Negative Aha	QUEST must be able to identify not only what it knows, but also what it does not know. This is termed the "known unknown." There may be a primitive non-semantic level of qualia where we can sense the object but we have been unable to generate an appropriate link set.
Intent: Theory of Mind (ToM)	Theory of mind is the act of computing the quale of "mindness" by an entity. It is one of the most important links for the quale of self. ToM allows for the understanding of others as intentional agents whose behavior is goal and perception driven. Mental states inside other entities are, by definition, unobservable. However, we can implement simulation techniques within our own mentally constructed world model to gauge another qualia agent's intent from observed activity. This tenet enables us to ascertain intent from observed activity.

Table 1 Some QUEST Tenets

you are stabbed. We are able *via* introspection to characterize that internal representation as having certain properties. The internal representations used in QUEST solutions all retain those characteristics. Our research effort has defined a set of 68 tenets, which at an abstract level describe our best understanding of how we should take a philosophical idea like qualia and physiological examples of what is possible provided by our own pattern recognition abilities and turn them into an engineering solution. Table 1 presents the most relevant to the N-QUEST domain. For a much deeper discussion of QUEST and the tenets, see Rogers' paper. [1]

Our primary inspiration for N-QUEST is the need to protect the USAF's unclassified computer network. This network consists of approximately 5 million devices, nearly all of which can be instrumented with security-oriented sensors. Gaining and maintaining situational awareness (SA) across a network this large and sensor-rich using the very best of today's technology are completely impractical. Something fundamentally new is needed. We propose that a hierarchical, qualia-based computing infrastructure will enable us to fully make use of all 5 million sensors in acquiring an understanding of the state of the network, its assets, and users, and making this understanding available to enhance the SA of the network warfare forces managing the network. Toward this end, we will strive to understand the qualia associated with patterns of activities across the network infrastructure and build a computing infrastructure capable of working with these qualia.

This work differs from current signature-based and anomaly-based intrusion detection in that we believe the qualia-based computing architecture may be able to not only truly "learn" low-level self and non-self patterns, but also develop higher level understanding of the relationships between activities and user behaviors, make predictions about its understanding of the meaning behind activity, and interact with the sensor fleet to verify or refute predictions. Although N-QUEST itself is novel, most of the components underlying its capabilities are drawn from the traditional body of work in this area. We believe current solutions will provide additional inputs into the N-QUEST representations and solutions.

Figure 2 illustrates our current understanding of how we may create a quale-based representational model, populate it with knowledge, and allow it to continuously sense and interact with the real world. This kernel represents a node in the overall self-awareness capability, and a single moment in time and space. Key components of the kernel include mechanisms in the traditional machine learning domain and sensory space. We anticipate using traditional pattern recognition techniques coupled with sensor feeds to provide inputs to a set of reasoning engines. These combine the sensed activity with knowledge and past experience to create and maintain a quale-based model of the real world and activities within it.

The model will contain not only understanding combined from knowledge and experience, but will also be able to accept sensor feeds and manage the sensor fleet to provide needed information. The underlying concept-based representational model includes the network itself; its links and nodes; its users and their behaviors; the missions that its users and itself support; and the behavioral interactions between related entities. This is not a simple model; the infrastructure itself is made up of millions of nodes and links, its users number in the hundreds of thousands, and the interaction between the users and infrastructure in performing the USAF's missions are myriad. As the representational models of the various entities in the kernel develop, sensory data is washed across them. Responses from the models in terms of "that feels familiar" similarities will enable the reasoning engines to predict the meaning of sensed activities. "Complete" matches result in AHA activities, whereas partial matches will be handled by managing the sensor fleet to update the world model in ways that support or refute the prediction.

An example of anticipated activity in the context of the initial fictional attack scenario might have the adversary using a compromised user account as part of the attack. The N-QUEST world model parts closest to the user would see logon and authentication from a known user. These parts would direct user-stimulated sensor flows across the myriad concepts embodied by that user's model at that moment, across N-QUEST's representation of the user's typical behavior. While a known user may have a representational node in the world model, the learned sense of that user will reside in the links between all parts of the model touched in any way by the user's activities. The true learning involved will include the links between model entities and also their patterns of activations as the entities go about their business. From this perspective, the kernel in Figure 2 represents a single quale and is tied to the other nodes through learned behaviors and links. Because the attack behavior does not match its sense of that user's "self," N-QUEST would predict that the user in question is not actually logged in and working. It would support that hypothesis by querying the building access control systems or by other methods such as interacting with the user to determine whether the user had entered the building. The combination of determining that the suspect user is actually an attacker and having the awareness of the user's activities across the network enables the system to control the user's ability to affect the overall domain. This information also enhances the SA of the network warfare forces.

The eyes, ears, and effectors of the network will likely be provided by the Cybercraft fleet. Cybercraft is an ongoing AFRL and AFIT research project that focuses on providing a trusted deployment platform for the USAF's network infrastructure. Essentially, the Cybercraft fleet will consist of hardened software and hardware-based mechanisms implanted throughout the network infrastructure. Each node will serve as a deployment platform for sensors, effectors, and decision engines. Between the nodes and controlling them is a hardened command and control network and infrastructure. [4]

Cyber Warfare Forces

To create effects in and through cyberspace, the Air Force must first enable combatant commanders to gain and maintain cyber superiority. In line with the national and military objectives, the Air Force has identified the following objectives:

- Deter and prevent cyberspace attacks against vital US interests
- Rapidly respond to attacks and reconstitute networks
- Integrate cyber power into the full range of global and theater effects
- Defeat adversaries operating through cyberspace
- Provide freedom of action in cyberspace for US and Allied commanders
- Create persistent cyberspace SA.

The technical capabilities embodied by N-QUEST are only part of the solution. We also need to develop a warfighting force capable of achieving the above cited objectives. In the USAF, we are doing so by modifying the mission and capabilities of the traditional communications and information (C&I) community, the electronic warfare community, and the

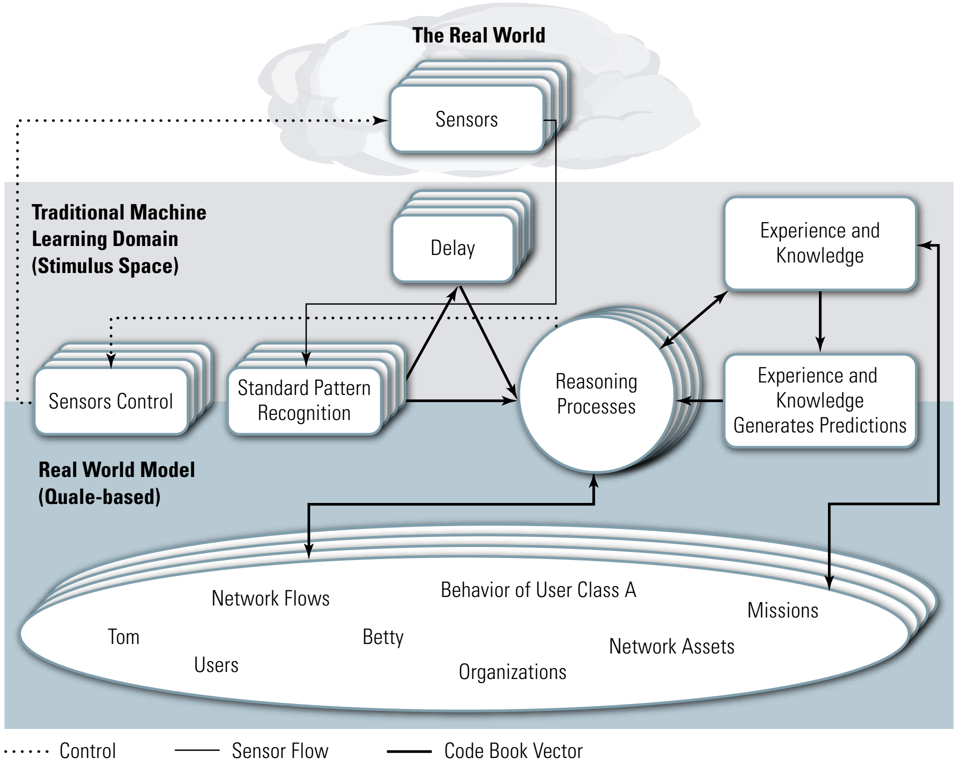

Figure 2 A Single Node in the Network QUEST Kernel

intelligence community. For the C&I community the focus must change from support to operations. Complete development of the cyber professional includes a dramatic cultural transformation. The growth of the cyber profession must take us from today's cyber culture characterized as "supporting" or "enabling" to a culture exemplified as "warfighting" and "operational." Cyber capabilities provided by today's electronic attack mission set bring a well-developed warfighting culture that must be matured in directed energy and network warfare operations; in the command, control, and operations of the USAF's cyber enterprise; and in the global provisioning, protection, and sustainment of cyber capabilities. This culture must be developed in our forces and operations across the vast cyberspace enterprise (wired and wireless capabilities; open- and closed-network operations; Internet Protocol (IP) and non IP-based connectivity; and terrestrial, airborne, and space-borne links and nets).

The USAF, under the guidance of Headquarters Air Force, A3O-C, is working to create a new cyber career force. Some parts of this force will provide the traditional support services in terms of networking, but the officers and some of the enlisted force will focus primarily on holding the adversary at risk and preserving our own freedom of maneuver in the cyberspace domain, and their specialty and expertise will stem from attack and defense skills. These forces will be capable of using mechanisms such as N-QUEST in flying and fighting those segments of cyberspace under their protection.

The Real Problem

The science fiction in the first part of this article, and the science and force development in the later sections actually address a relatively minute problem—the USAF's 5 million node network is just a small part of the DoD network and an even smaller part of the critical information infrastructure that the United States relies on every day. An attack against the United States will almost certainly span the entire military, government, and civilian infrastructure. We must explore the legal, policy, organizational, and technology issues surrounding this eventuality and prepare to defend our entire information infrastructure.

We hope this article will spur thought and discussion about ways to solve some of these urgent needs. The 5 million node problem is not manageable now, yet we must figure out ways to handle the larger issues! The high-risk research we are undertaking may not result in success; however, we are certain that we will learn much that will lead us to where we need to go to prevent "Tomorrow Night" from happening without the ability to respond. ■

References

1. Rogers, S. K., C. Sadowski, K.W. Bauer, M.E. Oxley, M. Kabrisky, A.S. Rogers, and S.D. Mott, "The life and death of ATR/Sensor Fusion and the hope for resurrection," Proceedings of SPIE, Volume 6967, March 2008.
2. Lacey, Tim, Robert Mills, Richard Raines, Paul Williams, and Steve Rogers,"A Qualia Framework for Awareness in Cyberspace," *IAnewsletter*, Vol. 10, No. 3, Fall 2007, pp. 12-17.
3. Ahn, L. von, M. Blum, N.J. Hopper, and J. Langford, CAPTCHA Web page; *http://www.captcha.net.*
4. Stevens, Michael and Paul D. Williams, "Use of Trust Vectors for Cybercraft and the Limits of Usable Data History for Trust Vectors," IEEE Symposium, Proceedings of the Computational Intelligence in Security and Defense Applications, CISDA 2007, Honolulu, HI, April 2007, pp. 193-200.

About the Authors

Major Paul Williams, USAF, PhD | earned his BS from the University of Washington, his MS from the Air Force Institute ofTechnology, and his PhD from Purdue University. He is an Assistant Professor of Computer Science and Cyber Operations in the Department of Electrical and Computer Engineering at the Air Force Institute of Technology, Wright-Patterson AFB, Ohio. He has served in many information operations roles, both operational and supporting, for seventeen years. His research interests center on cyber warfare, and include algorithms, artificial intelligence, and security-focused computer architectures.

Dr. Steve "Capt Amerika" Rogers | is a Senior Scientist at the Air Force Research Laboratory where he serves as the principle scientific authority for Automatic Target Recognition and Sensor Fusion. Dr. Rogers' research focuses on qualia exploitation of sensor technology, QUEST. Prior to coming back to work for the government Dr. Rogers founded a company for developing practical applications of advanced information processing techniques for medical products. Among the products developed by the company was the world's most accurate computer aided detection system for breast cancer.

Dr. Robert "Bob" Mills | is an assistant professor of Electrical Engineering in the Department of Electrical and Computer Engineering at AFIT. Dr. Mills received a BS degree in Electrical Engineering with highest honors from Montana State University, an MS degree in Electrical Engineering from AFIT, and a PhD in Electrical Engineering from the University of Kansas. His research interests include digital and spread spectrum communications; low-probability-of-intercept and anti-jam communications and networks; signal detection and exploitation; and mobile communication networks and security.

Dr. Richard "Rick" Raines | is the Director of the Center for Information Security Education and Research (CISER) at AFIT. Dr. Raines received a BS degree in Electrical Engineering from the Florida State University, an MS degree in Computer Engineering from AFIT, and a PhD in Electrical Engineering from Virginia Polytechnic Institute and State University. He teaches and conducts research in information security and global communications.

Mr. Timothy "Tim" Lacey | is an instructor of Computer Science in the Department of Electrical and Computer Engineering at AFIT. Mr. Lacey received a BS in Computer Science/Management of Computer Information Systems (magna cum laude) from Park College and an MS in Computer Systems from AFIT. He has several professional certifications to include Certified Information Systems Security Professional (CISSP), Certified Ethical Hacker (CEH), and Microsoft Certified Systems Engineer (MCSE). He teaches and conducts research in both network and software security.

Electronic Voting Security

by Angela Orebaugh

With the 2008 US Presidential election just around the corner, memories of past election controversies are certain to be on the minds of voters. The 2000 US Presidential election resulted in controversy over the winner of Florida's votes, ultimately leading to a recount and the Bush versus Gore Supreme Court case. The 2000 US Presidential election controversy resulted in the passing of the Help America Vote Act (HAVA) in 2002. The goals of HAVA are as follows—

- Replace punch card voting systems
- Create an Election Assistance Commission (EAC) to assist in the administration of federal elections
- Establish minimum election administration standards.

HAVA provided government funding to replace punch card and mechanical lever voting systems with new systems. Many states were allocated funding to upgrade their voting systems to new electronic systems manufactured by several different vendors.

In the 2004 US Presidential election, roughly 50 million votes were cast using electronic voting (e-voting) machines, whereas 32 million votes were cast with punch cards. [1] The 2004 US Presidential election concluded with numerous concerns such as inaccessible voting and vote counting inaccuracies. One main issue that arose was the accuracy and reliability of electronic voting machines.

E-voting machines include various computerized methods, such as electronic means of casting and counting votes. These methods may include punch cards, optical scan systems, and direct-recording electronic (DRE) voting systems. E-voting may also include the transmission of votes *via* communication lines (*e.g.*, land or cell phones, computer networks, and the Internet).

E-voting machines offer several advantages over traditional paper ballot voting methods such as improved vote counting speed and disabled voter accessibility.

E-voting machines offer several advantages over traditional paper ballot voting methods such as improved vote counting speed and disabled voter accessibility. Electronic ballots save costs by eliminating the need to print paper ballots and mitigate the possibility of running out of ballots during an election. However, e-voting machines may also include potential flaws and weaknesses and facilitate electoral fraud. Several electronic voting systems were identified in numerous reports as containing significant vulnerabilities. For example, Princeton University demonstrated that someone with physical access to a voting machine or its removable memory card could install malicious code that could steal votes undetectably and modify all records, logs, and counters. They also demonstrated an ability to create malicious code that spreads

Types of E-Voting Devices

- **Punch Card Machines**—These first generation electronic voting machines have been used since the 1960s. Voters punch holes in the ballot cards next to their choice, and electronic tabulation methods read the punches to total the votes.
- **Optical Scan Voting Systems**—Voters select their choices by making a mark directly on the ballot; marks are read using an electronic tabulation method that total the votes.
- **Direct-Recording (DRE) Voting Systems**—These systems use an electronic display to present a ballot and provide a computerized method for collecting and tabulating all votes in a single voting machine. DREs also may transmit vote totals to a central processing location and print a verifiable record for the voter.
- **Hybrid E-voting Machines**—These machines have electronic ballot-marking devices such as touch screens or other systems with electronic means for recording or tabulating votes.
- **Internet Voting**—Voting may be performed in traditional voting locations with Internet connected voting systems or from any Internet capable computer.

UNUM
01PLEASE SELECT A CANDIDATE01010
001PLEASE SELECT A CANDIDATE010
001PLEASE SELECT A CANDIDATE01010

…in a 2006 election in Florida, some votes intended for Democratic candidates were displaying as having been cast for Republican candidates. This problem was the result of calibration errors in the touch screen of the voting system.

automatically and undetectably from machine to machine, creating a voting machine virus. [2] Another possible attack involves an attacker physically inserting a hardware device on the voting machine to manipulate recorded votes. [3]

Many problems with e-voting machines have been identified in actual elections. For example, in a 2006 election in Florida, some votes intended for Democratic candidates were displaying as having been cast for Republican candidates. This problem was the result of calibration errors in the touch screen of the voting system. [4] In 2004, Montgomery County, MD, faced myriad problems with its new touch screen e-voting systems, including failure to boot up, screen freezes, smart card and encoder problems, and unexplained error messages. [5] In 2003, Fairfax County, VA, also experienced problems with its new touch screen e-voting machines, ranging from casting votes to reporting results. [6]

Several technologies (*e.g.*, paper verification and cryptography) exist that may be used to detect possible fraud or malfunction, audit the voting machine, and assure voters that their vote was cast correctly. A Voter Verified Paper Audit Trail (VVPAT) is often used to allow the voter to visually verify his or her choices before casting the vote. The VVPAT is often treated as an official ballot of record. E-voting systems may use cryptographic methods with mathematical calculations to enable voters to verify their vote is recorded and tabulated correctly. These systems often include an electronic receipt signed with a digital signature; however, these receipts must also guarantee voter anonymity.

Review and testing procedures may detect fraudulent code or hardware. One method to test voting machines is parallel testing, which compares an independent set of results with the original machine results. Parallel testing involves removing a randomly selected voting machine from service and testing it with voting test ballots. [7] Logic and accuracy testing (L&A) uses test votes during pre-election testing to determine if voting machines are functioning properly. Another method of testing is independent software verification and certification, which ensures the integrity of electronic voting machines by certifying and signing code. The code can then be verified to ensure it has not been changed before or during an election. However, many security experts recommend that voting software be open to public scrutiny.

The EAC maintains federal responsibility for accrediting voting system test laboratories and certifying voting equipment through the Voting System Certification and Laboratory Accreditation Program. This program independently verifies that voting systems comply with the functional capabilities,

Vulnerabilities of E-Voting Systems

- **Electronic Voting Offers Opportunities and Presents Challenges**
  http://www.gao.gov/new.items/d04766t.pdf
- **Federal Efforts to Improve Security and Reliability of Electronic Voting Systems are Under Way, but Key Activities Need to be Completed**
  http://www.gao.gov/new.items/d05956.pdf
- **RABA Trusted Agent Report for the State of Maryland**
  http://www.raba.com/press/TA_Report_AccuVote.pdf
- **SAIC Report**
  http://bravenewballot.org/resources/SAIC.pdf
- **Analysis of an Electronic Voting System**
  http://avirubin.com/vote/analysis/index.html
- **The Machinery of Democracy: Protecting Elections in an Electronic World**
  http://www.brennancenter.org/dynamic/subpages/download_file_39288.pdf

accessibility, and security requirements necessary to ensure the integrity and reliability of voting system operation, as established in the Voluntary Voting System Guidelines (VVSG). The National Institute of Standards and Technology (NIST) helps the EAC with the accreditation program through its National Voluntary Laboratory Accreditation Program (NVLAP) by providing recommendations to the EAC regarding laboratory accreditation. [8] NIST also helped create the 2007 VVSG guidelines, which were released in draft form in August 2007. [9] The VVSG provides a set of guidelines and requirements to increase the security, usability, and accessibility of e-voting systems. These voluntary guidelines contain requirements for vendors when developing voting systems and for laboratories when testing whether the systems conform to, or meet, the requirements of the guidelines. NIST is currently not only assisting the EAC with comments on the 2007 VVSG from the public review period but also developing an open, comprehensive set of test suites that will enable test laboratories to uniformly and consistently test voting systems against the 2007 VVSG requirements. [10]

Informed voters are an important defense against potential attacks. If you are casting your vote on an e-voting system with a VVPAT, please check your VVPAT before casting your vote. The larger the number of voters who check their VVPAT before casting their vote, the less likely that an Automatic Routine Audit will be unable to identify a Trojan horse attack.

The Information Assurance Technology Analysis Center (IATAC) has been instrumental in developing the 2007 VVSG and associated test suites. IATAC subject matter experts work directly with NIST computer scientists to write security guidelines and develop derived testing requirements for e-voting systems.

As you venture into the polls on 4 November 2008 to cast your vote for the next President of the United States, what can you do to help in the security of e-voting? Informed voters are an important defense against potential attacks. If you are casting your vote on an e-voting system with a VVPAT, please check your VVPAT before casting your vote. The larger the number of voters who check their VVPAT before casting their vote, the less likely that an Automatic Routine Audit will be unable to identify a Trojan horse attack. If you are using a paper ballot that an optical scan e-voting system will be reading, mark your choices carefully in accordance with instructions. The more voters who fill out their ballots correctly, the less likely that a Trojan horse attack on the over/undervote protection or scanner calibration will affect numerous recorded votes. [11] Lastly, immediately report any errors or suspected malfunctioning of an e-voting system to a poll worker. ■

References

1. New Study Shows 50 Million Voter Will Use Electronic Voting Systems, 32 Million Still With Punch Cards in 2004. *http://www.edssurvey.com/images/File/VotingEquipStudies%20/ve2004_news.pdf*
2. Security Analysis of the Diebold AccuVote-TS Voting Machine. *http://itpolicy.princeton.edu/voting*
3. Nedap/Groenendaal ES3B Voting Computer, a Security Analysis. *http://www.wijvertrouwenstemcomputersniet.nl/images/9/91/Es3b-en.pdf*
4. "Test Run for Voting" Miami Herald, 10/31/2006
5. Emerging Scandal on MD Voting Machine Performance. *http://www.truevotemd.org/Press_releases/html/2005-03-08_Press_Release.html*
6. Fairfax to Probe Voting Machines. *http://www.washingtonpost.com/wp-dyn/articles/A54432-2003Nov17.html*
7. *http://www.cs.uiowa.edu/~jones/voting/testing.shtml#parallel*
8. EAC's Testing and Certification Program for Voting Systems. *http://www.eac.gov/voting%20systems/docs/faqs.pdf/attachment_download/file*
9. Next Version Voluntary Voting Ssytem Guidelines (VVSG). *http://vote.nist.gov/vvsg-report.htm*
10. "National Institute of Standards and Technology's Role in Voluntary Voting System Guidelines and Testing." *http://vote.nist.gov/speeches/PublicMtg-100407-Testimony-Skall.htm*
11. The Machinery of Democracy: Protecting Elections in an Electronic World. *http://www.brennancenter.org/dynamic/subpages/download_file_39288.pdf*

About the Author

Angela Orebaugh | supports a variety of security engagements with the National Institute of Standards and Technology (NIST). She has 15 years experience in information technology and security and is the author of several technical security books including Nmap in the Enterprise and Wireshark & Ethereal Network Protocol Analyzer Toolkit. Ms. Orebaugh is also an adjunct professor at George Mason University. She may be reached at *iatac@dtic.mil*.

University of Virginia Security SMEs

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 SMEs profiled in this article are the University of Virginia's (UVA) IT Security and Policy Office administration.

Shirley Payne is Director for IT Security and Policy at the University of Virginia. In this capacity she focuses on the continuous enhancement of information technology policies and security of the university's diverse and decentralized computing environment. She works in partnership with units and individuals across the university to formulate policies, assess security risk, establish strategic direction, provide security education and training, implement security safeguards, track security incidents, develop business continuity plans, and related activities. Ms. Payne participates as part of the Vice President's senior staff in short and long term planning and budgeting, resolution of department issues or concerns, development of information technology standards and guidelines, and communications between the Vice President and the range of constituencies served by the ITC department. She tracks security and policy trends, issues, and best practices in these areas and keeps ITC, university senior management, and others informed. Ms. Payne has many years of experience in information technology, most of which has been in higher education. She holds a Bachelor's degree in Computer Science from Winthrop University and a Master's degree in Management Information Systems from the University of Virginia. Ms. Payne contributes to state and national level security and policy initiatives through workshop participation, presentations, and publications. She has recently presented on a number of information security topics including privacy and awareness. Ms. Payne is a member of the EDUCAUSE 2009 Program Committee and the Council on Technology Services (COTS).

Brian Davis is an IT Security and Policy Specialist for the IT Security and Policy Office. He led the design and implementation of a university wide IT Security Risk Management Program (ITS-RM). He is project manager for the University's Social Security Number Remediation Initiative. Mr. Davis also performs security awareness and training, security incident response, policy development and implementation, security and policy collaboration with the Virginia public higher education community, and legislative advising to the university on IT issues. Mr. Davis received his BA and MA from Emory University. He is a frequent speaker at EDUCAUSE, Who's Watching Charlottesville, and other IT security events.

Marty Peterman is an IT Security Specialist for the IT Security and Policy Office. He spearheads security awareness initiatives for the university and the greater Charlottesville area. Mr. Peterman leads the "Who's Watching Charlottesville?" cyber security initiative. He has presented at a number of security conferences, most recently at the EDUCAUSE Security Professionals Conference 2008 on "Community Aware: Taking cyber Security Awareness to the Street". He received his BA and MEd from Shippensburg University. Mr. Peterman currently holds GCWN, GCIH, and CISSP certifications.

If you have technical questions for a member of the IT Security and Policy office at the University of Virginia or another IATAC SME, please contact *http://iatac.dtic.mil/iatac*. The IATAC staff will assist you in reaching the SME best suited to helping you solve the challenge at hand. If you have any questions about the SME program or are interested in joining the SME database and providing technical support to others in your domain of expertise, please contact *iatac@dtic.mil*, and the URL for the SME application will be sent to you. ■

Recent Developments in Cyberlaw

by Rick Aldrich

Editor's note: This article does not constitute legal advice but is intended merely to raise awareness about key legal issues in cyberspace. Always consult your agency's attorney for legal advice.

Changes in technology and computer usage have resulted in several interesting new developments in cyberlaw. This article highlights only a few of those developments. The US border appears to be one of the newest frontiers for new cases. As hard drives have increased in capacity, even as computers have grown smaller, people increasingly have been carrying with them larger quantities of data than ever before.

Border agents have determined this trend to present a potential treasure trove. As such, they have been taking advantage of the border search exception to the Fourth Amendment to search laptops and other electronic devices of international travelers. The Fourth Amendment protects against "unreasonable searches and seizures," [1] but the Supreme Court has long recognized border searches as reasonable and therefore has not imposed the requirement of a search warrant on such searches. It has done so "pursuant to the longstanding right of the sovereign to protect itself by stopping and examining persons and property crossing into this country." [2] The border search exception also has been applied to international airports, where air travelers first touch US soil after crossing the border. Generally, border searches can be performed without any suspicion whatsoever. Even for "nonroutine" border searches that are especially intrusive (*e.g.,* strip searches or body cavity searches), the courts have required only a "reasonable suspicion," far less than the "probable cause" standard applicable to searches requiring a search warrant.

The Fourth Amendment protects against "unreasonable searches and seizures," [1] but the Supreme Court has long recognized border searches as reasonable and therefore has not imposed the requirement of a search warrant on such searches.

Should courts consider a search of a computer at the border as a routine search requiring no reasonable suspicion or a nonroutine search requiring reasonable suspicion? Most courts have held that such searches are routine; [3] however, one court recently determined such a search to be nonroutine and therefore suppressed the results of the search. [4]

In the *Arnold* case, an individual flew nearly 20 hours from the Philippines to Los Angeles International Airport. Upon his arrival, Customs agents requested that Arnold turn on his computer to demonstrate that it worked and subsequently noted two folders named "Kodak Pictures" and "Kodak Memories." A cursory inspection found some images that caused Customs agent to seize his computer. A couple of weeks later, agents found images of child pornography. At Arnold's trial, he moved to suppress the images as the product of an unlawful search. The United States defended the search as lawful under the border search exception. Federal district court judge Dean Pregerson held that the intrusiveness of the search made it a nonroutine search, which should have been justified by reasonable suspicion, but finding none, suppressed the fruits of the search. He held that—

> *[w]hile not physically intrusive as in the case of a strip or body cavity search, the search of one's private and valuable*

personal information stored on a hard drive or other electronic storage device can be just as much, if not more, of an intrusion into the dignity and privacy interests of a person.

This presents a potential conflict with *United States v. Romm,* [5] another case out of the Ninth Circuit, in which the search of a traveler's computer at the border was deemed to be covered by the routine search provision of the border exception. The Ninth Circuit heard the Government's appeal of the *Arnold* case on 15 October 2007 and is expected to issue a decision in the case soon. Some court observers have opined that the Ninth Circuit seemed skeptical of Judge Pregerson's rationale and were inclined to overrule it. Whatever the court's decision, it has sparked controversy among many who travel and wonder what their rights are when the government seeks to search or seize a computer at the airport that might have sensitive information on it, such as proprietary information, intellectual property, privileged communications, or other important information that they cannot do without for any extended period of time. Indeed, this has even prompted a lawsuit by the Electronic Frontier Foundation and the Asian Law Caucus. [6]

The *Romm* case also raised another interesting issue relating to what constitutes "knowing possession." Romm was charged with knowing possession of child pornography. He alleged that he viewed the images only at a website, but never saved them to his hard drive so as not to "possess" them. What Romm apparently did not realize is that the act of viewing the images on his computer caused the images to be automatically saved to his computer's cache. [7] The question presented to the court was whether one could be convicted of a crime requiring "knowing possession" if the only evidence of the possession came from the cache. In *Romm,* the court held it was sufficient. The court seemed persuaded by expert testimony indicating that knowledgeable users could access the image in cache at will, save it under another file name, and otherwise exercise traditional notions of possession.

Not all courts have decided this issue in the same manner. Indeed, in *United States v. Kuchinski,* [8] the court refused to consider 19,000 child pornography images found in the defendant's computer cache as a valid basis for increasing his sentence under the Federal Sentencing Guidelines. The court held,—

Where a defendant lacks knowledge about the cache files, and concomitantly lacks access to and control over those files, it is not proper to charge him with possession and control of the child pornography images located in those files, without some other indication of dominion and control over the images. [9]

As such, it appears the courts will continue to wrestle with this issue.

Another intriguing border case raised issues relating to encryption. As concerns have grown over privacy and data protection issues, businesses, governments, and users have increasingly turned to encryption as one answer. Sebastian Boucher, a Vermont resident, encrypted his Z: drive, on which he stored images that he downloaded from the Internet, including pornographic images. When he crossed from Canada back into the United States with his computer in December 2006, border agents asked him if he had any child pornography on his computer. With amazing frankness, he indicated he was unsure. He explained that he downloaded adult porn and that sometimes child porn came mixed in with it, that he tried to delete it, but was unsure whether he had deleted all of it. The border agents asked him to show them the drive. Unbeknownst to them, Boucher entered a passphrase to unlock his Z: drive in response to their request. When the agents saw images they believed to be child porn, they seized his computer and then turned it off.

Two weeks later, when the agents tried to access the images they had seen, they realized they could no longer access the drive because it was encrypted. The agents obtained a subpoena, ordering Boucher to provide the passphrase. He refused, arguing that ordering him to divulge his passphrase would violate his right against self-incrimination under the Fifth Amendment. [10] The Government then sought to compel Boucher to type in his passphrase in private without saying anything and then turning the computer back over to the Government. Boucher refused, again invoking the Fifth Amendment. The court held such a procedure was still "testimonial" and so still violative of the protection against self-incrimination. [11] The result would likely have been different if Boucher had written down the passphrase because then a search warrant, or even a subpoena, could have been used to obtain it. But because Boucher held the passphrase only in his mind, the court held that the Government could not force him to reveal it.

The Government would be permitted to attempt to crack the encryption, but a strong encryption key would render a brute force attack a nearly futile effort. As users move toward encrypting an increasing amount of data stored in files or encrypting entire drives, this case could pose new challenges for law enforcement, counterintelligence agents, and system administrators.

Another interesting case dealt with encryption and the scope of third-party consent. In *United States v. Andrus*, [12] federal authorities came to suspect Ray Andrus of involvement in child pornography, but because they lacked probable cause for a search warrant, chose to employ a "knock-and-talk" approach. This approach involves knocking at the door of the suspect and trying to engage the suspect in conversation with the hope that something the suspect says will provide a sufficient basis for probable cause.

The authorities went to the place where Ray Andrus lived with his 91-year old father, Dr. Andrus. The elder Andrus answered the door in his pajamas. After asking some preliminary questions to validate that Ray lived there, the authorities asked if Dr. Andrus would consent to a search of the computer that his son used. Dr. Andrus consented and led them to his son's room. The agents began a forensics search of Ray Andrus's computer and promptly found evidence of child pornography. Later, the agents determined that the images on the computer were password protected. Under fairly well-established law, a third-party can consent to a search only to those things over which the person shares equal access and control with the owner. Because agents belatedly learned that Dr. Andrus did not usually enter his 51-year old son's room, did not use his son's computer, and did not know the password to enter his son's password-protected account, Ray challenged the evidence as illegally obtained, without valid consent. He lost at the trial level, and subsequently on appeal.

Forensics software can easily circumvent the password protection of a user account, but when one is operating under third-party consent, one is restricted by the authority of the third party.

The appellate court reasoned that because Dr. Andrus lived with his adult son, apparently paid the Internet bill, was able to access his son's room through an open door, and the computer could be viewed from outside the room, Dr. Andrus had "apparent authority" on which the agents could reasonably rely. As for the password protection, the court held that because the agents did not realize the account was password protected before their search, the court would uphold the seizure of the evidence. This case is troubling because the court seems to suggest that agents who ask few questions and proceed blindly will be given more latitude than those who proceed meticulously.

Forensics software can easily circumvent the password protection of a user account, but when one is operating under third-party consent, one is restricted by the authority of the third party. Here, the court seems to suggest that quickly conducting a forensic search before one learns of limiting facts will be rewarded.

An interesting trend in computer usage relates to the increasing tendency of users to store their data remotely. Remote backup storage, remote tax filing storage, and remote storage of calendars are only some examples of this trend. But entrusting data with a third party has legal consequences, which translates to a reduced expectation of privacy because of the risk that the third party could choose to share your data with law enforcement.

This issue was at the crux of the case of *Warshak v. United States.* [13]

The Government had secured an ex parte order requiring Warshak's email provider to disclose the contents of his "old" [14] email, pursuant to a statute authorizing such. [15] The order required his Internet service provider (ISP) not to inform him that his emails were being disclosed to the Government and the ISP complied.

The Government was required to inform Warshak within 90 days, but for reasons unclear from the case, did not alert him for a year. Once Warshak learned of the order, he sued to bar the Government from obtaining any further such orders on the basis that he had a reasonable expectation of privacy in all his emails, whether "old" or not, and that the disclosure to the Government therefore violated his Fourth Amendment protection against unreasonable searches and seizures.

The district court agreed, so the Government appealed. On appeal, the Government argued that Warshak could not have a reasonable expectation of privacy in emails he stored with a third party, his ISP, because the ISP had the right under law to read his email under certain conditions. The court held it was not as important what the third party *could* do as what in practice they *did* do. Because the ISP rarely read individual's emails, the

▷▷ *continued on page 30*

University of Virginia, IT Security and Policy Office

by Angela Orebaugh

The University of Virginia (UVA), founded in 1819 by Thomas Jefferson, is nestled in the shadows of the Blue Ridge Mountains in Charlottesville, VA. UVA has nationally acclaimed business, medical, and law schools and is home to over 20,000 undergraduate, graduate, and professional students each year. The IT Security and Policy Office is responsible for coordinating, developing, and enforcing computer security and policies across UVA's diverse and decentralized computing environment. The IT Security and Policy Office reports directly to the Vice President and Chief Information Officer (VP/CIO) [1] and is part of the Information Technology and Communication (ITC) department. [2]

The IT Security and Policy Office works in partnership with units and individuals across the university to establish strategic directions, provide security education and training, assess security risks, implement security safeguards, detect and respond to security incidents, develop business continuity/disaster recovery plans, comply with federal and state regulations, and related security activities. It advises senior executives and managers on security issues and risks. The office also works with advisory committees, legal counsel, and others to develop, implement, and keep current a comprehensive set of policies governing the university's information technology resources. It provides ongoing policy interpretation, education, and advice to the University community, and it works with various University units and law enforcement to effectively address policy or law violations. As legislative advisor to the UVA State Governmental Relations Office, the office provides assessments, advice, and general information on proposed legislation regarding information technology issues. It also has responsibility for reviewing and commenting on new or changed state IT policies, standards and procedures, responding to state requests for IT information, facilitating university compliance with state IT reporting requirements, and keeping appropriate university officials apprised of state IT issues and concerns. Additionally, the Director of IT Security and Policy serves on the State's Council on Technology Services, Information Security Council, and chairs the Virginia Alliance for Secure Computing and Networking (VA SCAN). [3]

In an ongoing effort to secure UVA's technology infrastructure and core services, ITC has commenced a $1.2 million hardening and securing program. This program will advance the University's goal of having an information technology infrastructure that has a level of redundancy and resistance to threats that is appropriate for the university. The program consists of three areas—

- Securing sensitive data
- Eliminating single points of failure in mission-critical systems and services
- Implementing a set of tools for stress testing systems and applications.

To secure sensitive data, the university has issued a university-wide Social Security Number (SSN) policy that includes security awareness, data classification, data security standards, data stewardship, tools for identifying high sensitivity data inventories, and data security remediation plans and implementation. To eliminate single points of failure the university will focus on upgrading network storage appliances for email and other core services, replicating and load balancing networks, relocating web clusters, upgrading power, and performing space and cooling assessments. Next fiscal year's hardening work will leverage the storage infrastructure built this year and will include services such as ITC's server virtualization environment, Exchange, and other high-priority UNIX- and Windows-based services. The university also plans to implement tools to stress test the first phase of the new Student System. [4]

UVA is also a founding member of the Virginia Alliance for Secure Computing and Networking (VA SCAN). VA SCAN's purpose is to strengthen information technology security programs within the Commonwealth of Virginia. The Alliance brings together Virginia higher education security practitioners, who developed and maintain security programs widely emulated by other institutions, and researchers responsible for creating cyber

▷▷ continued on page 28

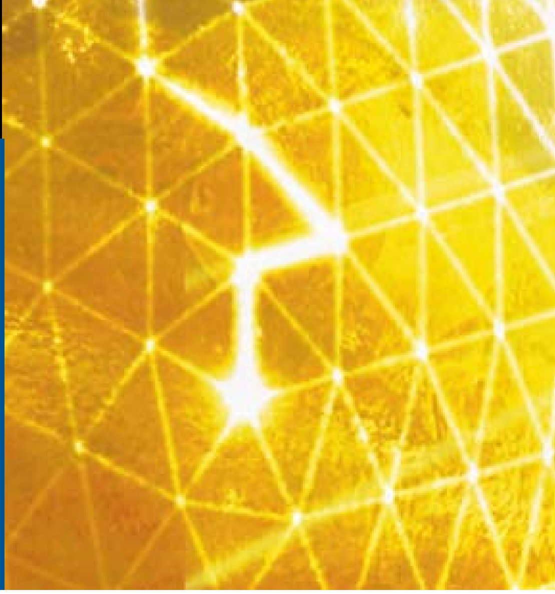

Securing the Converged Enterprise, Part I

by AT&T

Convergence Trends

In the context of computer and telecommunications networks, "convergence" has historically meant combining voice, data, and video on a common network. However, convergence is now acquiring several other definitions.

For example, many employees work from home. They share personal computers (PC), local area networks (LAN), and Internet access connections with family members, using the same physical circuits to connect to the public Internet and their employer's private intranet. Public and private network traffic is converging on commercial digital subscriber line (DSL), cable modem, and other broadband access lines.

Meanwhile, wired and wireless networks are merging to support consistent application experiences as users roam. This phenomenon, commonly known as fixed-mobile convergence (FMC), is in its infancy. Dual-mode smart phones and other devices that support connections to mobile wide area networks (WAN) and wireless LANs (WLAN) are part of the FMC landscape.

Applications that allow single user identities and phone numbers to work across the entire blended wired/wireless infrastructure are now available. Presence capabilities, which combine user location and availability information for managing personal communications, are being introduced to help stitch applications into a seamless experience.

Soon, specialized networks will plug into the corporate network as another form of network convergence. Among these are radio frequency identification (RFID) networks (for asset tracking and supply chain management), sensor networks (for remotely monitoring and controlling industrial devices), and closed-circuit television (CCTV) video surveillance networks (for blending physical security with information technology [IT] resources). These specialized networks will connect to the enterprise's traditional LAN *via* Internet Protocol (IP) and web services, providing enterprise-wide access.

The merging of networks, traffic types, applications, and interfaces makes life simpler and more productive for end users. From a cost perspective, capital and operational expenses required for running one network, rather than several isolated "silos," drop significantly. Merged networks and applications also open up new opportunities for correlating data and events across the organization. This ability improves employee decisionmaking and enhances customer service in call centers and elsewhere throughout the business.

An ability to correlate network events—specifically, security-related events—can actually enhance the ability to protect the converged enterprise network. On the other hand, convergence introduces some new security risks that the IT department must address.
Those risks and some general advice for mitigation are discussed in this article, which introduces the following recommendations for securing converged enterprise networks—

- A "defense-in-depth" approach to security, which employs multiple layers of user screening and encryption
- Centralized management of security helps deliver improved levels of security and scalability efficiencies
- Integration of security components into network devices, which simplifies the security infrastructure and renders it less likely to fail.

At Issue: Complexity Increases Risk

Enterprise IT departments must balance benefits of convergence with associated new security risks. Although convergence eases communications and data access tasks for the typical end user, complexities associated with supporting multiple interfaces, protocols, and devices create scalability challenges for IT that can lead to potential vulnerabilities.

Voice, data, and video, for example, might traverse any number of access networks because telecommuters, road warriors, and extranet partners now use a wide variety of devices and interfaces to connect to back-end resources. A given organization may support hundreds of wired or wireless interfaces to a public WAN, virtual private network (VPN) service, cellular network, or public switched telephone network (PSTN). Collectively, these interfaces represent a large, complex

set of network entry points that the IT department must manage and protect.

Software Update Challenges and Vulnerabilities

Organizations have traditionally prioritized the sequence in which various sites receive security and operating system (OS) patches and upgrades. They base the priority patching on the criticality level associated with each site. For example, the data center usually takes top priority.

Software patches are now released frequently, and the sheer number of sites and interfaces to be updated is proliferating quickly. An IT staff may not get beyond updating the first few priority locations before returning to the site at the top of the list to apply still newer patches. At some point, a hacker could conduct a distributed attack to identify a weak point in the infrastructure, such as a lower priority site with out-of-date software. A worm or Trojan horse could then be introduced to the network that could affect availability and uptime. Similarly, severely out-of-date firewalls could provide an outsider with access to private files and databases, which the hacker could copy to steal data.

A centralized, automated system for issuing patch updates is very useful in combating this issue, and this system will be discussed in the section, Centralized Security Management.

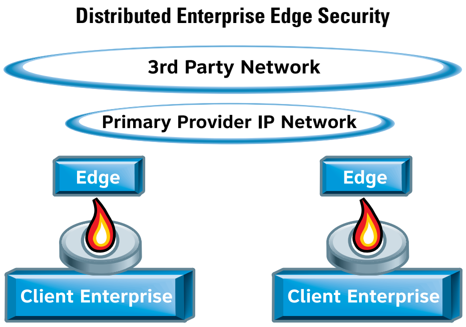
Distributed Enterprise Edge Security
3rd Party Network
Primary Provider IP Network
Edge
Edge
Client Enterprise
Client Enterprise

Premises-Based Solutions

- Major business security investment at edge
- IDS, Firewalls, Anti-Virus, Anti-SPAM deployed by business
- Broad-based network attacks difficult to defend against at individual locations
- Disparate security policies for Internet connected endpoints
- Difficult to scale

Figure 1 This diagram depicts a Premises-based Security Solution.

Growing Internal Threats

Adding to the complexity of risk mitigation is that internal attacks also have become a growing issue. An increasing number of security attacks come from inside an organization. For example, the Computer Security Institute's (CSI) 2006 Computer Crime and Security Survey report revealed that 7 percent of several hundred respondents attributed more than 80 percent of cyber crime losses in 2005 to insiders. [1] This means that a single firewall sitting at the perimeter of the network between LAN and WAN, although still necessary, may no longer be adequate to secure the entire converged enterprise.

Governance Mandates

Finally, organizations now must comply with the latest corporate and industry governance mandates (*e.g.,* Sarbanes-Oxley, Gramm-Leach-Bliley [GLB], Health Insurance Portability and Accountability Act (HIPAA), Basel II, and Payment Card Industry Data Security Standard). These mandates require log-in audit trails and resource access tracking.

Creating a defense-in-depth network security infrastructure that protects against intrusions at multiple network segments can greatly assist in keeping all sites and interfaces updated and protected, internally and externally. It also can help ensure compliance with the security components of government mandates.

Defense-in-Depth

A defense-in-depth approach to network security creates a network infrastructure that is highly resilient to internal and external attacks. Building defense-in-depth security entails deploying different forms of security in various places throughout the converged network to mitigate a mix of risks. By setting up multiple checkpoints between a user attempting to gain network access and the intended destination data resource, organizations can accurately verify user access rights. Checkpoints also can be effective at scanning network traffic for malicious code that might disrupt service. The code can then be filtered off the network.

Network junctions where one type of network connects to another, where users must cross from one LAN segment to another, or devices in which user groups and departmental traffic are logically separated can be considered network trust boundaries. For defense in depth, user authentication checks and scanning for malicious code should occur at the primary trust boundaries—

- Between public and private network interfaces using network firewall and intrusion detection system (IDS)/intrusion prevention system (IPS) products and services
- Between LAN segments and internal departments using firewalls and IPSs
- At the "mobile edge" in client devices with endpoint security software and encryption/VPN software.

Protection at these key network junctures establishes a resilient underlying LAN-WAN platform that protects multiple types of application traffic. Traditional convergence, the merging of voice, data, and video onto a common infrastructure, opens the door to a single attack being able to potentially affect all these types of application traffic. Reinforcing the base infrastructure has become extremely important.

Centralized Security Management

By centralizing the management of these defense-in-depth security components, an organization can achieve the scale needed to help ensure higher levels of security. Taking a network-centric security approach means creating one central place for setting, maintaining, and enforcing a common set of security policies across all network sites. This setup enables businesses to overcome the patch-vulnerability issue discussed in Section I. By pushing software updates out to predetermined network devices simultaneously from a central location, organizations avoid the one-at-a-time update approach that can result in some sites having outdated software that is vulnerable to attack.

Premises-Based Solutions

Premises-based solutions are as follows—

- Major business security investment at edge
- IDS, firewalls, antivirus, anti-SPAM deployed by business
- Broad-based network attacks difficult to defend against at individual locations
- Disparate security policies for Internet-connected endpoints
- Difficult to scale.

The centralized location could be a service provider's security operations center. In this case, businesses would subscribe to a carrier's security service. This effort involves pushing out updates to all sites and devices per individual corporate policy while using the provider's economies of scale. Depending on the network segment to be protected (*e.g.,* LAN or WAN), security appliances may or may not be needed on the customer premises.

Let's examine the basics of securing the various network segments that comprise a defense-in-depth architecture.

Securing the WAN

In the WAN, network-based security, often in the form of a service, uses a series of gateways in the service provider's network that reside between users and data resources. The gateways translate private IP addresses into publicly routable addresses. This system helps ensure that a private device never directly exposes its IP address to the public Internet, PSTN, or other shared network. This prevents a hacker from piggybacking onto that address for entry into private network resources or launching another type of attack into the private network. The more gateways used, the deeper a hacker will have to penetrate to identify a private routable IP address. Using multiple gateways makes it more difficult for a hacker to succeed.

Similarly, network-based firewall services protect connections made between two WANs. Capable today of deep packet inspection, today's firewalls permit and deny access based on user access control lists (ACL). They also can filter anomalous signatures and protocol behavior as packets travel between networks, serving an IPS function. It is prudent for firewall-based filtering to occur between any two dissimilar networks, particularly between a public Internet service and a VPN (or private network). For example, if a business site or employee's home office uses a DSL or other commercial Internet access connection to reach its corporate multiprotocol label switching (MPLS) VPN service, filtering should occur where the access network meets the MPLS network. Another appropriate spot for filtering is between two corporate partner networks that both run Internet-based VPNs but allow some resource sharing between their networks.

If a services-based approach to network-centric security management is taken, providers also may offer additional scanning services and reports. Such services might scan the public Internet to detect precursors to worms and other events and send notification alerts of pending vulnerabilities. Services also may be customized by examining individual Internet or VPN traffic and potentially detecting a distributed denial-of-service (DDoS) aimed at that network. To mitigate risks, some managed VPN services will automatically deploy policies and take action when certain events are detected on the VPN.

For protecting the privacy of data in transit, encrypted VPNs should be used when traffic traverses the public Internet infrastructure. Encryption scrambles data and authentication information. To protect the privacy of data in transit, encryption can be used to create a private "tunnel" for each customer through the publicly shared Internet. The VPN can be in the form of an IP Security (IPSec) VPN service between fixed corporate sites or a secure sockets layer (SSL) VPN service for remote and mobile users.

Many enterprises using MPLS VPN services elect not to encrypt traffic because MPLS technology creates virtual circuits that keep the customer's traffic from intermingling with that of others. However, companies with the highest security requirements, such as financial institutions transmitting customer account data, may elect to encrypt their MPLS traffic as an approach for double security protection. Services are available to encrypt the traffic across the shared MPLS backbone network segment.

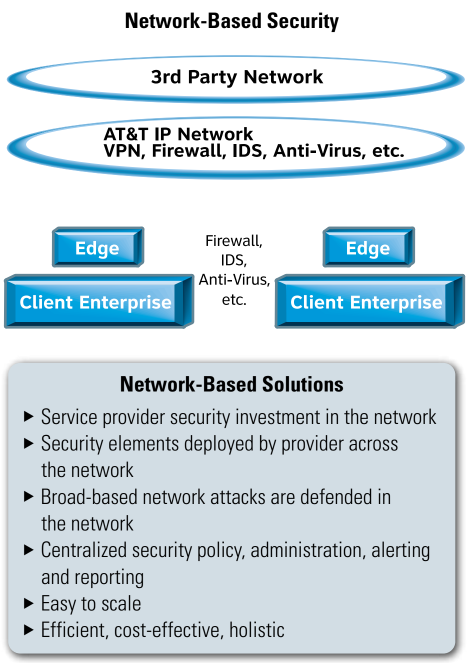
Network-Based Security
3rd Party Network
AT&T IP Network
VPN, Firewall, IDS, Anti-Virus, etc.
Edge
Client Enterprise
Firewall, IDS, Anti-Virus, etc.
Edge
Client Enterprise
Network-Based Solutions
▸ Service provider security investment in the network
▸ Security elements deployed by provider across the network
▸ Broad-based network attacks are defended in the network
▸ Centralized security policy, administration, alerting and reporting
▸ Easy to scale
▸ Efficient, cost-effective, holistic

Figure 2 This diagram depicts security built into a network, protecting business network and applications.

Securing the LAN

Given that incidents of internal attacks are increasing, security between LAN segments and between LAN application servers (places that represent internal trust boundaries) has become another priority. For example, IPSs focus on filtering anomalous or otherwise suspicious traffic off the LAN at internal trust boundaries. When managed centrally, an operations center would continually send updates with the latest known malicious signatures to IPS appliances that sit between the access network and distribution network (wiring closet) and between the distribution network and core LAN in the data center.

Network-Based Solutions

- Service provider security investment in the network
- Security elements that the provider deploys across the network
- Broad-based network attacks defended in the network
- Centralized security policy, administration, alerting, and reporting
- Easy to scale
- Efficient, cost-effective, and holistic

Securing the VoIP Network

To a large degree, protecting the Voice over IP (VoIP) network involves the same set of protective services that have long been in place for data network infrastructures. If VoIP (and IP video) are simply new applications being added to the IP network, it is difficult to secure that traffic if it is running over a vulnerable infrastructure.

Just as data might be separated into virtual LANs (VLAN) for different user groups, with different resource access rights belonging to different VLANs, voice traffic may be segregated onto its own VLAN. This helps ensure that VoIP devices can talk only to other VoIP equipment and cannot use the VoIP network as a launching pad into the data network. There also is a quality-of-service (QoS) benefit to putting VoIP on its own VLAN, which can be prioritized for low latency.

In some ways, VoIP is simpler to secure than it was in the traditional circuit-switched environment. For example, encryption of voice calls for privacy is possible in the packet-switching environment, where this was not previously available. Most VoIP vendors encrypt in the handsets they sell, so conversations are protected end-to-end.

Guarding against toll fraud, or theft of service, involves the same basic practice as in circuit switching. Here, extension transfers to outbound ports are disabled. For off-LAN calls, using multiple gateways between an IP address and the public Internet (*i.e.,* PSTN) (depending on where a call is terminating) prevents the handset's private IP address from being exposed as a possible attack point.

Securing Wireless Networks

WLANs, also called 802.11 and wireless fidelity (Wi-Fi) networks, have inherent authentication and encryption for use over the LAN. To help ensure privacy and avoid theft of user credentials in public Wi-Fi hotpots, many organizations rely on IPSec VPNs to encrypt over-the-air data when Wi-Fi client devices are used remotely.

On the corporate campus, a possibility exists that unauthorized or rogue devices might associate to the network. Similarly, a personal wireless client device might erroneously associate to a rogue 802.11 radio. If malicious, it might attempt to grab user credentials (a breach called wireless phishing). Thwarting attempts to steal credentials involves deploying the latest version of 802.11 authentication and encryption standards. Preventing rogue radios from flooding Wi-Fi client devices with bogus disassociation messages, thus overloading the device and causing denial of service, requires radio frequency (RF) specific IPSs. These IPSs are often sold as a third-party overlay system or service or might be bundled into a basic WLAN system.

Securing Remote and Mobile Endpoints

Endpoint security plays a key role in mobile networking. It is important to keep endpoints (or clients) free from viruses and other malware and to remain in compliance with corporate standard software versions for OSs, security, and application software. Taking a centralized, network-centric approach to endpoint security keeps policies and security versions consistent throughout the converged enterprise.

If a device is lost or stolen, encrypting the hard drive of endpoint devices storing mission-critical data will help protect against data theft. Personal firewalls running on the endpoints block hacker intrusions into the device for data theft or for piggybacking onto a corporate network connection.

Application convergence, the blending of mobile and wired networks, and the telecommuting phenomenon are creating new requirements for securing endpoints, or client devices. One requirement is to prevent devices from passing infected code to the corporate network. While traveling, a user might unplug from the corporate network, connect to the public Internet, and pick up a virus or other malware. Businesses must guard against viruses impacting the user's local data and prevent viruses from being transmitted to the corporate network.

Conclusion

The traditional WAN perimeter is still vulnerable and continues to require firewall-based protection. However, security in the converged enterprise no longer represents protecting only one physical network perimeter. Instead, there are now multiple network "edges," requiring a distributed, defense-in-depth security architecture with WAN gateway services, firewalling, IPS, and endpoint security.

The first step in securing a converged network is to ensure that the underlying infrastructure is reinforced with these capabilities. A centralized, network-centric approach will provide added protection layers by automatically handling, deploying, and maintaining the latest versions of security system software, OSs, and application software. Managed security services also can add another layer of protection by scanning WAN traffic, alerting network customers about detected events, and possibly taking automated actions when events are identified.

Ensuring all sites and interfaces are continually in compliance with security, OS, and application software versions will shut down the occasional open network pinhole. This effort will help prevent distributed attacks that could exploit the point of entry, preventing data theft or the introduction of malware onto the network.

Part II of this paper examines the various security components and services in greater depth to offer more detailed understanding of the role each "layer" of security throughout the enterprise. ■

References

1. Lawrence A. Gordon and Martin P. Loeb, Computer Security Institute," 2006 CSI/FBI Computer Crime and Security Survey," page 12, figure 13.

For additional information about this topic, visit http://www.att.com/networkingexchange or call Mary Forbes 443/896-5221

▷ *continued from page 23*

IATAC SPOTLIGHT ON EDUCATION: UVA

security instruction and research programs nationally recognized for excellence. The goals of the alliance are to—

- Help avoid costs associated with security breaches
- Save security program development time
- Reduce security training costs
- Take advantage of economies of scale.

VA SCAN offers services such as policy examples and templates, self-assessment checklists, training and awareness materials, and expert advice. VA SCAN won the 2005 Award for Excellence in Information Technology Solutions from EDUCAUSE. [5]

UVA is known for its community involvement and computer security is no exception. UVA is a partner in the "Who's Watching Charlottesville?" [6] initiative that aims to promote cyber security awareness in the Charlottesville area. UVA was recently recognized by the Association for Computing Machinery's (ACM) Special Interest Group for University and College Computing Services (SIGUCCS), a national association for higher education IT professionals, for its short video, "The Job Interview," [7] which won a "Best of Category" Communications Award. The video was part of ITC's contribution during National Cyber Security Awareness Month and addresses the consequences of personal information on the Internet. ■

References

1. More information on the VP/CIO Office of Information Technology can be found at *http://www.itc.virginia.edu/oit/org/home.html*
2. More information on the Information Technology and Communication department can be found at *http://www.itc.virginia.edu/org*
3. More information on the IT Security and Policy Office can be found at *http://www.itc.virginia.edu/org/security*
4. *http://www.itc.virginia.edu/projects/hardening/update040708.html*
5. *http://vascan.org/news/index.html*
6. *http://www.whoswatchingcharlottesville.org*
7. *http://www.whoswatchingcharlottesville.com/videos/TheJobInterview.mov*

More Focus Required on Web Applications

by Allan Carey

Organizations demand that more features and capabilities be available through the applications leveraged to provide services on the web. More applications are becoming web-enabled and exposed through intranets and the Internet.

Businesses, governments, and other organizations have been aggressively pushing network boundaries to achieve interoperability and interconnectivity with their business partners and customers. The efficiencies, opportunities, and services derived from extending business processes and applications outside the traditional controllable network environment to the Internet have not come without a price. Increased risk continues to be front and center consistently for executives and senior leadership.

Organizations demand that more features and capabilities be available through the applications leveraged to provide services on the web. More applications are becoming web-enabled and exposed through intranets and the Internet. To increase information sharing, collaboration, productivity, and the customer experience, new and more complicated coding methods are being employed, which in turn increases the complexity to manage, patch, and update those applications. Some applications older than 10 years are considered to be too fragile to update or patch, yet they are being accessed from the web.

The methods used by, and constraints placed on, application development teams have not changed dramatically. Often, software is still pushed into general availability without adequate testing because the software needs to be on time and within budget. Throughout their careers, developers have acquired neither proper training nor a level of awareness pertaining to security or information assurance (IA). Only now is the subject and its associated importance coming to light.

Web application security is critical to organizations in private and public sectors for numerous reasons (*e.g.,* regulatory compliance such as Payment Card Industry Data Security Standards (PCI), the need for business resiliency, and new risks being introduced through a larger attack surface). Any organization handling credit card data is required to perform an annual application pen test under PCI. Organizations are more dependent than ever on information technology (IT) and therefore must have resiliency built into their computing infrastructure. The attack surface is expanding exponentially as applications become web enabled and as new technologies such as Web 2.0, blogs, and social networking sites are leveraged for conducting business.

Taking into account the latest research, the problem is worsening. In September 2007, Symantec released its Threat Report, which stated that 61 percent of all vulnerabilities reported were related to web applications. From the SANS Top 20 Internet Security Risks, number one of the top new risks that are particularly difficult to defend was critical vulnerabilities in web applications. Based on a recently published report by the Web Application Security Consortium (WASC), its web hacking incident database reported that more than 44 percent of 2007 incidents were tied to noncommercial websites such as Government and Education. Overall, the number of incidents reported grew from 44 in 2006 to 82 in 2007.

To mitigate threats from web applications, the process must start at the core of the organization. Some steps to consider are as follows—

- Determine the risk relevance of various web applications by threat modeling parts of the business
- Define a clear accountability model for the development and maintenance of the web-facing applications
- Evaluate the return on investment based on the results of the threat modeling exercise
- Where applicable, benchmark yourself against other organizations from a process, effectiveness, and investment perspective
- Build a strategy to address the problem, including a budget
- Socialize the strategy with the application development leader(s) to discuss time and resource allocation requirements. Typically, little to no coordination exists between IA and the application development groups
- Most importantly, gain senior leadership buy-in to champion and drive the web application security initiative. ■

References

1. *http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_internet_security_threat_report_xii_09_2007.en-us.pdf*
2. *http://www.sans.org/top20/2007/press_release.php*
3. *http://www.webappsec.org/projects/whid/statistics.shtml*

About the Author

Allan Carey | is the Senior Vice President of Research and Product Development at the Institute for Applied Network Security (IANS). In this position, he manages all research and intellectual property across the Institute. Prior to IANS, Mr. Carey spent seven years at IDC, a global provider of market intelligence and advisory services for the IT sector. He developed and managed the Security Services practice and provided in-depth analysis, intelligence and consulting on key aspects of the information security and business continuity services markets. He may be reached at the Institute for Applied Network Security, 15 Court Square, Suite 1100, Boston, MA 02108, by telephone at 617/399-8100, or by email at *acarey@ianetsec.com.*

▷ *continued from page 22*

RECENT DEVELOPMENTS IN CYBERLAW

court held that Warshak retained a reasonable expectation of privacy.

This is an interesting distinction that the courts did not recognize previously and may require government service providers to change their procedures. It appears that if they want to negate an expectation of privacy in their users, they are best advised not only to reserve the right to monitor electronic communications (generally through banners and user agreements) but also to *actually* review them periodically to meet this "actual practice" test.

As technology and the ways in which people use computers continue to evolve, so will the law evolve. This evolution might occur in fits and spurts occasionally and will sometimes involve conflicting opinions along the way, but for this reason it is especially important that information assurance professionals and investigators stay abreast of the developments. ■

References

1. The Fourth Amendment to the US constitution states in full: The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.
2. *United States v. Ramsey*, 431 US 606, 616 (1977).
3. See, *e.g., United States v. Ickes*, 393 F.3d 501 (4th Cir. 2005).
4. *United States v. Arnold*, 454 F. Supp. 2d 999, 1003-04 (N.D. Cal. 2006), appeal docketed, No. 06-50581 (9th Cir. 13 June 2007).
5. 455 F.3d 990 (9th Cir. 24 July 2006).
6. "Suit: Airport searches of laptops, other devices intrusive," available at *http://www.cnn.com/2008/TRAVEL/02/11/laptop.searches*.
7. A cache is "a temporary storage area where frequently accessed data can be stored for rapid access." *Wikipedia.org.* Frequently, web browsers will store images in cache to reduce the time to download the page on future visits.
8. 469 F.3d 853 (9th Cir. 2006).
9. *Kuchinski, supra* note 9.
10. The Fifth Amendment to the US constitution states, in pertinent part, "No person shall be … compelled in any criminal case to be a witness against himself."
11. *In re Boucher,* 2007 WL 4246473 (D. Vt., 29 November 2007)
12. No. 06-3094 (10th Cir., 25 April 2007).
13. No. 06-4092, 2007 WL 1730094 (6th Cir. June 18, 2007).
14. Refers to email that has been in storage for more than 180 days, under 18 U.S.C. 2703(a).
15. 18 U.S.C. 2703(b).

About the Author

Rick Aldrich | is the IATAC Senior Computer Network Operations Policy Analyst. Previously, he served as the Deputy Staff Judge Advocate for the Air Force Office of Special Investigations, specializing in the cybercrime and information operations portfolios. He has multiple publications on cyberlaw issues and has been a speaker on the topic at numerous national and international conferences. Mr. Aldrich holds a BS in Computer Science from the US Air Force Academy, a JD from UCLA, and an LLM in Intellectual Property Law from the University of Houston.

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 ________________________________

Organization ________________________________

Address ________________________________

____

____

DTIC User Code ________________

Ofc. Symbol ________________

Phone ________________

Email ________________

Fax ________________

Please check one:

□ USA
□ Industry
□ USMC
□ Academia
□ USN
□ Government
□ USAF
□ Other
□ DoD

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

____

LIMITED DISTRIBUTION

IA Tools Reports (softcopy only)

□ Firewalls
□ Intrusion Detection
□ Vulnerability Analysis

Critical Review and Technology Assessment (CR/TA) Reports

□ Biometrics (soft copy only)
□ Data Mining (soft copy only)
□ Wireless Wide Area Network (WWAN) Security
□ Computer Forensics* (soft copy only. DTIC user code MUST be supplied before these reports will be shipped)
□ Configuration Management
□ IA Metrics (soft copy only)
□ Defense in Depth (soft copy only)
□ Network Centric Warfare (soft copy only)
□ Exploring Biotechnology (soft copy only)

State-of-the-Art Reports (SOARs)

□ Data Embedding for IA (soft copy only)
□ Modeling & Simulation for IA (soft copy only)
□ Software Security Assurance
□ IO/IA Visualization Technologies (soft copy only)
□ Malicious Code (soft copy only)
□ A Comprehensive Review of Common Needs and Capability Gaps
□ The Insider Threat to Information Systems

UNLIMITED DISTRIBUTION

IAnewsletters Hardcopies are available to order. The list below represents current stock.
Softcopy back issues are available for download at *http://iac.dtic.mil/iatac/IA_newsletter.html*

Volumes 4		□ No. 2	□ No. 3	□ No. 4
Volumes 5	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 6	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 7	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 8	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 9	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 10	□ No. 1	□ No. 2	□ No. 3	□ No. 4
Volumes 11	□ No. 1	□ No. 2		

Fax completed form to IATAC at 703/984-0773

Calendar

July

FISC 2008
1–2 July 2008
Colorado Springs, CO
http://www.fbcinc.com/fisc/

53rd Joint Eletronic Warfare Conference
8–9 July 2008
Lackland AFB, TX
http://www.fbcinc.com/event.aspx?eventid=Q6UJ9A00FJ5F

SERP 2008
14–17 July 2008
Las Vegas, NV
http://www.world-academy-of-science.org/worldcomp08/ws/conferences/serp08/General%20Information

WorldComp 2008
14–17 July 2008
Las Vegas, NV
http://www.world-academy-of-science.org/worldcomp08/ws

August

Black Hat
2–7 August 2008
Las Vegas, NV
http://www.disa.mil/conferences/index.html

LinuxWorld Conference and Expo
4–7 August 2008
San Francisco, CA
http://www.linuxworldexpo.com/live/12/

DEFCON 16
8–10 August 2008
Las Vegas, NV
http://www.defcon.org/

Crypto 2008
17–21 August 2008
Santa Barbara, CA
http://www.iacr.org/conferences/crypto2008/

DoE IT Day
20 August 2008
Washington, DC
http://www.fbcinc.com/event.aspx?eventid=Q6UJ9A00GA27

September

Network Security Conference
8–10 September 2008
Las Vegas, NV
http://www.isaca.org/Template.cfm?Section=Network_Security_Conference&CONTENTID=38203&TEMPLATE=/ContentManagement/ContentDisplay.cfm

IT Security World 2008
15–17 September 2008
San Francsico, CA
http://www.misti.com/default.asp?page=65&Return=70&ProductID=7154

NATO Standardization Conference
16 –17 September 2008
Lansdowne, VA
http://www.fbcinc.com/event.aspx?eventid=Q6UJ9A00G0XZ

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