Volume 13, Issue 2 April, 2014

ISSN 1445-3312 (Printed Journal)

ISSN 1445-3347 (On Line Journal)

JOURNAL OF INFORMATION WARFARE

Volume 13 Issue 2 April 2014

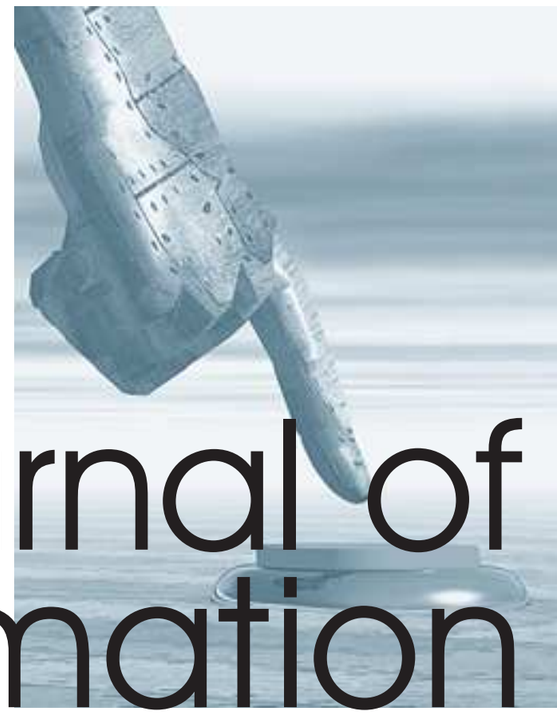

Journal of Information Warfare

Journal of Information Warfare (JIW)

www.Jinfowar.com

Scope

The journal has been created to provide a forum for discussion, information, and interaction between practitioners and academics in the broad discipline of information warfare/operations. It is of interest to professionals from the military, government, commerce, industry, education, and academy.

A full gambit of topics is covered, from physical destruction of information systems to the psychological aspects of information use. The aim is to provide a definitive publication that makes available the latest thinking and research in the critical area of information warfare.

The Journal of Information Warfare is published four times per year and is available both online and in hard copy.

Authors' Responsibilities & Copyright

Authors are to ensure the accuracy of their papers. This journal does not accept any responsibility for statements made by authors in their written papers. Where relevant, authors are to ensure that the contents of their papers are cleared for publication, for example, by their employer, their client, the funding organization and/or copyright owner of any material that is reproduced.

The views expressed in contributions to this Journal do not necessarily represent those of the editors, Advisory Board, or the publishers.

Subscription

Individual, Student, and Corporate subscriptions are available. For current pricing, see http://www.jinforwar.com/subscribe/.

Individual
This is a twelve-month subscription to the journal for individual subscribers. This is a download version only. Hardcopies can be purchased if required**.**

Individual, Student
This is a twelve-month subscription to the journal for students. Evidence of full-time study must be provided. This is a download version only. Hardcopies can be purchased if required

Corporate
This is a twelve- month subscription to the journal for corporate/library subscribers. This includes a download version and a hardcopy when available. A single subscription covers unlimited use for a single campus/geographic location. Additional hardcopies can be purchased if required

Note: Hardcopy purchase is only available to subscribers.

All advertisements in this journal are printed free of charge as a service to readers.

Journal Cover Design, Concept and Layout
by Laima Croft

Journal of Information Warfare

Volume 13, Issue 2

Contents

Journal of Information Warfare

Published by:
School of Computer and Security Science,
Edith Cowan University, Western Australia
ISSN 1445-3312

Online version published by:
Mindsystems Pty Ltd
ISSN 1445 3347

http://www.Jinfowar.com

Editorial

To all readers, we are very excited about this issue. This is a special edition of *the Journal of Information Warfare* (*JIW*) and the first of its kind where we have collaborated with the Information Assurance Directorate (IAD) of the National Security Agency (NSA). In this publication, we bring you 10 articles from current and highly technical subject matter experts from NSA, all of which focus on cyber-security efforts that attempt to realize their theme of Confidence in Cyberspace. We hope you enjoy this special issue, and it is our desire to continue this new effort as an annual tradition.

As many of you know, the staff of *JIW* have been very involved and active in a number of IO and IW conferences as will be shown below. We recently attended and led several panel discussions at the latest International Conference on Cyber Warfare and Security (ICCWS '14) at Purdue University on 24-25 March in West Lafayette, Indiana. It was a great event and kudos to Sam Liles and Eugene Spafford for their leadership over the last year in making this happen.

There are also a number of other great IO and IW conferences that you should be aware of, all of which are also outstanding opportunities for you to meet and collaborate with other cyber academics:

- Information Operations Global, 17-19 June 2014, The Kensington Close Hotel, London, http://www.informationoperationsevent.com

- Spend the 4th of July in the birthplace of Democracy!!!! Go to Athens, Greece for the next European Conference on Cyber Warfare and Security, http://academic-conferences.org/eccws/ECCWS-home.htm

- The Australians also host a great series of Cyber Warfare conferences every year in the early December timeframe, normally in Perth, but it also rotates. Check out the URL, http://conferences.secau.org/

- Go on a Safari next year!!! The ICCWS will be heading to South Africa in the world famous Kruger National Park on 24-25 March 2015, http://academic-conferences.org/iccws/iccws-home.htm

Finally the staff of the JIW is endeavouring in every way to increase the usability of this publication. We have stood up the Cyber Base, which is a new site on which to share key information for *JIW* subscribers: www.mindsystemscyberbase.com.

Our goal is to work with you to expand this capability as we move forward. Cheers

Dr. Leigh Armistead, CISSP, CDFE
Chief Editor, *Journal of Information Warfare*
larmistead@gbpts.com

Forthcoming Information Warfare Conferences

Information Operations Global
The Kensington Close Hotel
17-19 June 2014
London, UK
http://www.informationoperationsevent.com

13th European Conference on Cyber Warfare and Security (ECCWS-2014)
University of Piraeus
3-4 July 2014
Athens, Greece
http://academic-conferences.org/eccws/eccws2014/eccws14-home.htm

2014 SRI Security Congress
Edith Cowan University
1-3 December 2014
Perth, Western Australia
http://conferences.secau.org/

10th International Conference on Cyber Warfare & Security
Kruger National Park
24-25 March 2015
South Africa
http://academic-conferences.org/iccws/iccws-home.htm

Authors

Vincent (Mike) Boyle is the Standards Lead for the Information Assurance Directorate (IAD) at the National Security Agency (NSA). He coordinates IAD's efforts in various public standards bodies and represents IAD in the Department of Defense standardization community. Mr. Boyle earned his bachelor's and master's degrees in Mathematics and Applied Mathematics from the University of Maryland Baltimore County, and has developed a strong background in cryptography and secure network protocols.

Patrick DeSantis is an analyst with the National Security Agency (NSA) Information Assurance Directorate (IAD). He now specializes in research of industrial-control-systems' security vulnerabilities and exploits in support of IAD's effort to secure National Security Systems and the national critical infrastructure. Mr. DeSantis earned master's and bachelor's degrees in Management Information Systems from the University of South Florida and holds numerous professional certifications, including Offensive Security Certified Professional (OSCP) and Certified Information Systems Security Professional (CISSP). Prior to joining NSA, Mr. DeSantis served as a Ranger in the U.S. Army, taught college-level computer science courses, and conducted professional information-security vulnerability assessments and penetration tests.

Gregory W. Hale was a Program Director with the NSA, where he led a team comprised of government and contractor personnel in developing performance standards to accredit companies that provide Intrusion Detection, Incident Response, Vulnerability Assessment/Analysis, and Penetration Testing services that are designed to meet the growing cyber defense needs of the U.S. government. Mr. Hale has an M.B.A. in Business Management from Touro University International and an M.S. in Information Assurance from Colorado Technical University. He has more than 39 years of combined military and civilian experience in the fields of electro-mechanical engineering, telecommunications, computer science, information assurance, and organizational leadership and management. Mr. Hale retired in February 2014.

Michael Herring currently serves as the Technical Director for the National Security Agency Active Cyber Defense Initiative and is responsible for identifying candidate technologies for integration into holistic ACD solutions to defend U.S. Government networks. He also teaches graduate courses in Risk Analysis, Process Strategy, and Information Technology Strategy as a member of the Loyola University Maryland Information Systems and Operations Management Department. Mr. Herring is a graduate of Mississippi State University with a degree in Electrical Engineering, and of Loyola University Maryland with a master's in Business Administration.

Ronald M. Lenzner is Systems Engineer and Project Director with the National Security Agency directing operational aspects of NSA's National Security Cyber Assistance Program. Mr. Lenzner holds a master's of science in Management Information Systems, earned a graduate certificate as an Information Systems Analyst, and was awarded the title of *Master* by the National Security Agency's Technical Track Program. He is also an adjunct facility instructor with National Cryptologic School for digital communications and has over 37 years with the Department of Defense.

Diana Parr is a Cybersecurity Technical Leader for the National Security Agency (NSA). She supports the NSA Information Assurance mission as a cybersecurity advocate to senior leadership within NSA, as well as

diplomatic and intelligence agencies. She has worked on projects involving computer network attacks against government and diplomatic systems, cloud computing vulnerabilities, nuclear command and control protection, and other information assurance issues. She also served in technical management roles involving data center management, help desk operations, and software development, both while working for NSA's Technology Directorate and in a previous assignment with the Office of the Director of National Intelligence (ODNI). She is the recipient of two national intelligence awards for service and leadership in the intelligence community. Prior to joining the federal government, Ms. Parr worked in private industry for 20 years, including four years working for software companies in Silicon Valley. Ms. Parr has an M.A. in Strategic Security Studies from the National Defense University in Washington, D.C., with concentrations in International Security Studies and Homeland Security Strategy and Leadership. She also holds a bachelor of science degree in Information Systems Management from the University of Maryland Baltimore County (UMBC).

Debora A. Plunkett serves as the National Security Agency's Information Assurance Directorate where she leads efforts to protect and defend United States' national security systems. She is the focal point for cyber security, cryptography, and information systems security for our nation's critical systems. With over 28 years of government service, she has held leadership positions in both the signals intelligence and information assurance missions at NSA. She previously served as a Director on the National Security Council at the White House for the administrations of both President Clinton and President George W. Bush, where she helped shape national policy and programs in support of critical infrastructure protection and cyber security.

Sue A. (Sandi) Roddy is the Technical Director, Analysis and Data Fusion, IAD. She provides technical leadership and strategic direction when performing technical assessments of Information Technology (IT) products and systems. She also provides technical leadership and strategic direction to achieve automation of analysis of cyber data in support of Department of Defense (DoD) and Intelligence Community (IC) networks and systems. Ms. Roddy has a B.S. in Computer Science from the University of Maryland, University College, and an M.S. in Information Technology from the University of Maryland, University College. She is also certified as an Information Systems Security Professional (CISSP), Information Systems Security Engineering Professional (ISSEP).

Berl M. Thomas is currently the Technical Director (TD) for the Information Assurance (IA) Operations Deputy Directorate at the National Security Agency. As the TD, he is responsible for providing technical and operational leadership to the IA Operations leadership team, personnel, and missions including hunting, blue team, red team, COMSEC monitoring, technical security, 24/7 operations, capability development and deployment, and cryptographic key production. Mr. Thomas holds a B.S. in Computer Science from Mississippi State University and an M.S. Computer Science from John Hopkins University.

Jeffery Watkins has worked at the National Security Agency for 30 years where he is currently serving as the CSfC Communications Manager. He graduated Summa Cum Laude with a bachelor of science degree in Information Systems Management (University of Maryland University College, 2004). During his career at NSA, Mr. Watkins has developed and fielded numerous secure voice system interfaces, provided information systems security engineering support to the Combatant Commands/Services/Agencies and

managed the certification of cross-domain solutions to meet Warfighters' requirements. Additionally, Mr. Watkins served in the field as the Senior NSA Liaison to the Defense Information Systems Agency (DISA).

Keith Willett is the lead architect for Active Cyber Defense (ACD). Mr. Willett has a B.S. in Computer Science with Mathematics minor from Towson University (1984); an M.S. in Business and Information Systems from University of Baltimore (1986); an M.S. in Information Assurance from Norwich University (2005); and is currently working on a Ph.D. in Systems Engineering Security at Stevens Institute of Technology (~2017). Mr. Willett holds $(ISC)^2$ CISSP and ISSAP certifications and has over 30 years' commercial and government experience in technology and security as an educator, programmer, database administrator, operations manager, systems engineer, enterprise architect, and enterprise security architect. Mr. Willett is the co-author of two books *How to Achieve 27001 Certification* and the *Official $(ISC)^2$ Guide to the ISSMP CBK*; and sole-author of the book *Information Assurance Architecture* all published by Auerbach Publishing.

Neal Ziring is the Technical Director for the National Security Agency's Information Assurance Directorate (IAD), serving as the technical advisor to the IAD Director, Deputy Director; and other senior leadership. Mr. Ziring is responsible for setting the technical direction across the Information Assurance mission space. He tracks technical activities, promotes technical health of the staff, and acts as liaison to various industry, intelligence, academic, and government partners. As part of his role, he guides IAD's academic outreach program, acting as the technical liaison to several universities that are participants in the National Centers for Academic Excellence—Research (CAD-R) program. His personal expertise areas include router security, IPv6, VM-based secure systems, cloud computing, cross-domain information exchange, and mobile code security.


VISION
Confidence
in CyberSpace
MISSION
Protect
Information
Outmaneuver Cyber Adversaries
INFORMATION ASSURANCE
DIRECTORATE

Achieving Confidence in Cyberspace in an Ever-Changing Ecosystem

DA Plunkett

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *The Information Assurance Directorate (IAD) of the National Security Agency is charged with developing security solutions that protect and defend National Security Systems. This cannot be accomplished by NSA alone. Partnerships with industry, academia, U.S. and foreign government entities are critical to delivering solutions that will meet and defeat cyber challenges of today and tomorrow. A comprehensive approach and strategy are needed. Key components of this defensive strategy include collaboration, security automation, resiliency, and a robust, sustainable cyber workforce. While the United States Government (USG) will certainly benefit from the successful implementation of these strategies, ultimately the entire global ecosystem will be more strongly protected by the advanced level of secure systems and capabilities available, and through a greater awareness of the harm that sophisticated and determined adversaries can cause to the cyber ecosystem.*

Keywords*: Collaboration, Security Automation, Cyber Workforce, Resiliency, Commercial Solutions for Classified, Active Cyber Defense, Threat Landscape, Cloud and Big Data*

Introduction

Without a doubt, these are times of great technological change. It could be argued that this rapid technological evolution started with the introduction of personal computers in the mid-1970s. That revolution took computers out of government, academia, and labs and put them into the homes and hands of people around the globe. In 1980, there were fewer than a million personal computers in the United States. By 2002, the billionth personal computer was sold. By 2015, it is estimated that there will be more than two billion PCs in use (Worldometers 2014).

The next significant technological development was the Internet—a concept that made information available and connected around the globe. And while the early Internet was not very capable or available, today it provides exponential improvements to knowledge sharing, access, business functions, and collaboration. In 2010, there were approximately 12.5 billion devices connected to the Internet. Already the number of things connected to the Internet exceeds the number of people using the Internet; by 2015, it is estimated that there will be 25 billion devices connected to the Internet—more than three times the projected human population (Evans 2011).

The Internet was quickly followed by the cellular revolution—the technology that provides access to data anywhere, anytime. There were an astounding 1.75 billion mobile phones and tablets sold in 2012. In fact, mobile computing now accounts for 12 percent of Americans' media

ISSN 1445-3312 print/ISSN 1445-3347 online

consumption time, triple its share in 2009 (Digby 2014). This trend shows no sign of slowing down.

One result of the rapidly changing technology ecosystem is the immediate availability and accessibility of information. It is estimated that the digital world will grow by a factor of 300 from 2005 to 2020, expanding from 130 exabytes to a staggering 40,000 exabytes, or 40 trillion gigabytes (Gantz & Reinsel 2012). This data will not be buried in filing cabinets, but will be stored online in largely discoverable, usable forms. Technologies surrounding cloud and big data provide the ability to aggregate data, ask specific questions, and discover things about the data that could not have been learned as easily otherwise. The availability of all of this data online provides much capability but also constitutes a new exposure, a new opportunity for exploitation. At the same time, however, these technologies also afford a new avenue for defense—for detecting attacks and exploitations so that they can be mitigated.

Evolving-Threat Landscape

Globally, people benefit significantly from advanced and emerging technologies, but these technological advances bring incredible vulnerabilities and significantly increased risks. The reality is that most networks are interconnected—including those that host important government and critical infrastructure functions. Reliance on technology and the Internet for social, commercial, business, and national security functions has resulted in the exponential growth of cyber threats as more threat actors seek ways to gain access to personal and business data through all available networks.

Before 2001, threat actors concentrated their exploitation on collecting and intercepting information. In fact, many early hackers were in it for the status. Hacking is now a business line for some; in fact, it is fairly easy to hire a hacker through online message boards to exploit systems, data, and credentials. By 2007, threat actors continued to practice exploitation techniques but also evolved from simply ‘breaking in’ and stealing information (that is, intellectual property) to Distributed Denial of Service (DDOS) attacks. These DDOS attacks can be very costly and have the potential to cause users to lose faith in their ability to protect and defend that which legitimately belongs to the information owner.

Now, integrity attacks (data deletion and corruption) are beginning to emerge and multiply. These types of attacks are potentially devastating and may be difficult to detect and deter. Of late, an even more disturbing trend is surfacing in which bad actors have the ability to cause destruction to computer systems and their information. A recent exemplar for such extreme actions was in the 2012 cyber attack against Aramco, a state-owned Saudi Arabian oil company. This attack was aimed at disrupting production from the world’s largest exporter of crude. More than 30,000 computers were affected by this attack (Mahdi 2012).

Threats escalate as technology evolves, and the value flowing through the economy is increasingly reliant on these technologies. The challenge of identifying emerging risks and vulnerabilities that result from the introduction of new and updated technology to public and private network infrastructures continues. These security challenges are not rooted in the fact that technology changes and evolves too quickly. These security challenges exist because the appropriate security responses have not been successfully implemented at the same pace at which

the technology has changed (and continues to change). In essence, poorly secured or even unsecured technologies are now the mainstay for the world's businesses and its citizens' personal lives.

The Mission of the IAD

The Information Assurance Directorate of the National Security Agency (NSA) has the responsibility of ensuring the government is on the right path to position this nation to meet future cyber-security challenges. While the IAD's mission is to protect National Security Systems (NSS), the government's reliance on commercial products creates the motivation to be very interested in the security of commercial components. In fact, practically all of the United States' government systems rely on commercial hardware and software—or are connected to networks that do. In light of this reality, the IAD must determine how best to achieve an acceptable measure of confidence in cyberspace.

To be sure, the focus must be on providing assurance through risk management—maintaining confidence and certainty in the security of the nation's systems and networks even in light of the potential for bad actors to do bad things to existing infrastructures. The IAD is establishing practices and processes that, when implemented, will allow the conduct of missions using the latest technological advances while maintaining trust in the networks and in the integrity of the data all Americans rely on. Because security risks and vulnerabilities occur in both hardware and software platforms, the IAD must take a holistic approach when devising solutions.

In order to achieve confidence in cyberspace, government, industry, and academic communities must come together to develop a common understanding of how all pieces interconnect—both hardware and software. No one entity has insight into everything—not government, not industry, not academia. The focus on confidence is, at its core, about reliability and trustworthiness and, yes, about assurance. In light of inevitable challenges, confidence will allow all the communities to proceed with their businesses even in light of adversaries who are determined to interrupt their abilities to do so.

So what are some steps the IAD of the NSA is taking to achieve an acceptable level of confidence in cyberspace while faced with vastly emerging technology trends such as mobility and cloud computing? At a high level, the agency is focusing its efforts on four areas: collaboration, automation, cyber-workforce development and extension, and resiliency.

Collaboration

The first and most critical step to achieve confidence in cyberspace is robust partnering. NSA's Commercial Solutions for Classified (CSfC) initiative is a great example of the power of partnerships. CSfC is positively affecting the security standards of commercial products that are not only used for government purposes, but are available on the market for anyone to use. It is centered on the agency's work with industry to specify NSA's security requirements and ensure industry products can be seamlessly and securely integrated into NSA architectures. The advantages of this program will extend past the confines of government architectures by providing benefits to critical infrastructure and the rest of the private sector that uses the security-enhanced products or protection profiles.

This collaboration is not limited to external industry partnerships. IAD partners with academia, other government agencies, and with internal NSA partners. For instance, the NSA IA mission gains significant benefit from partnerships with NSA's Signals Intelligence Directorate, benefits that include gaining insights into the intentions of foreign adversaries, which allows for the rapid and responsive development of security solutions in time to make a difference.

Automation

The next step to achieve confidence in cyberspace goes hand in hand with collaboration. It is automation—specifically, security automation, which enables system security, enhanced collaboration, and real-time response to network attacks. This automation is a key component of the IAD's overarching Active Cyber Defense Strategy. Ultimately, the IAD needs an open standards way to package and share all security-related information. There are currently multiple efforts underway in this area, including SCAP—Security Automation Content Protocol *(NIST, SP888-126 Rev.2, Version 1.2)* (NIST 2014), IODEF—Incident Object Description Exchange Format (*IODEF 50070, Version 00)* (Danyliw, R, Meijer,J & Demchenk,Y 2007), STIX—Structured Threat Information eXpression /TAXII (*Trusted Automated eXchange of Indicator Information*) (STIX 2014), and others. Adoption of a set of standards would result in usable mechanisms for sharing actionable information. In fact, the IAD is currently partnering with industry and government to determine the appropriate standards needed to meet information-sharing interests. Working together, industry and government will be able to achieve a common adoption of public standards. The truth is that working together to adopt a set of common standards and achieve consensus on those standards will create the ability to more easily share actionable information in real time. This ability to share in real time is a priority. All the members of the IA community—government, industry, academia—need to collaborate to develop and identify more efficient ways to learn from each other's insights and benefit from each other's advancements, without eliminating competition.

A critical component of the current cyber-security strategy is having a defensible architecture (hardware/software); networks must be defensible before they can be defended. Also important is having measures in place to empower the defender to focus on the hard problems while allowing the networks to be more self-healing. Security Automation is an essential attribute of a defensible network. Security Automation is important in order to efficiently scale defenses to expand with the growth of network infrastructure. When implemented, it enables automated intrusion-detection across multiple and even diverse systems. NSA is currently working with NIST and DHS to develop common standards for security automation, but the government alone will not be able to successfully achieve automation without significant partnership and buy-in from industry.

Over the past 40 years, the NSA has pushed its software security practices to places no one ever envisioned, and these efforts must continue to push the boundaries and move toward a more automated cyber defense. Still, too much time and resources are spent managing 'network hygiene' problems: missing patches, poor configurations, weak passwords, and unenforced policies. These are known classes of problems, with known solutions. NSA needs to make its networks defensible, with a much greater focus on large-scale standardization and automation. Making this goal a priority will protect the enterprise through automated best practices and enable the sharing of findings in real time. This automation would eliminate much of the noise and enable the redirection of limited resources to the hardest problems.

In short, focusing on increased collaboration and the sharing of actionable information as defined by the recipient will allow the agency to learn from others and will make all networks safer in the years to come. Collaboration and security automation are both critical to achieve confidence in cyberspace. No one entity has the visibility, control, or scope to do this alone. The technologies in use are too diverse, and they change too quickly—both the systems and the human resources must be ready and available.

Cyber-workforce development and extension

The next step toward achieving confidence in cyberspace is continuing to build the cyber workforce. Future cyber-security professionals must be armed with the skills and knowledge they need to perform a whole spectrum of information-assurance and cyber-security functions. To further reduce vulnerabilities in the national information infrastructure, higher education in information assurance/cyber security must be promoted to produce a growing number of professionals with expertise in various disciplines. It is imperative that the nation focuses on creating a larger, more technically diverse cyber workforce to meet the evolving cyber-security challenges.

Those challenges highlight why the role of academia is critical. NSA considers the task of helping the nation secure the next generation of cyber professionals one of its highest priorities. It is the reason the agency established the Centers of Academic Excellence (CAE) in 1999 with seven colleges teaching an established information-assurance curriculum. The intent of the National Centers of Academic Excellence program is to provide a mechanism to allow colleges and universities to engage with NSA and each other in building strong curriculums that will benefit the nation.

Department of Homeland Security (DHS) leadership recognized the shared IA mission, and a Memorandum of Agreement (MOA) between NSA and DHS was signed on 9 February 2004. At that time, DHS joined in partnership with NSA to co-sponsor the CAE initiatives. NSA continues to work closely with DHS on policy, criteria, and general activities concerning the development of the information-assurance field at colleges and universities across the country.

Today, NSA and DHS now jointly sponsor the CAE IA program. There are currently 181 CAE institutions in 43 States, the District of Columbia, and the Commonwealth of Puerto Rico. These institutions include research institutions, 4-year undergraduate and graduate programs, and community colleges and technical schools.

In 2013, NSA and DHS developed a new CAE knowledge unit to reflect the evolving nature of the IA/Cyber Defense (IA/CD) field today. The requirements for the new CAE IA/CD program are in alignment with DHS' *Cyber skills task force report of fall* 2012. The revisions to the program allow institutions to apply for Focus Areas (FAs) which are in-depth courses of study that institutions can offer to attract students looking for a specific Cyber Defense discipline. The program will also allow students to make an informed decision when choosing a program of study. Industry and government will eventually benefit from the revised program because hiring managers will be able to focus recruiting efforts on students from academic institutions with FAs that meet mission needs. In short, mapping to FAs will foster and encourage further development of strong IA-focused education and research depth at U.S. institutions.

The new program was modeled after NSA's CAE Cyber Operations Program, whereby institutions offer a core curriculum made up of Knowledge Units (KU), covering foundational IA topics such as Networking, IA Fundamentals, and Cyber Threats. CAE IA/CD institutions will also be able to identify a Focus Area within the curriculum, such as Secure Cloud Computing, Cyber Investigations, or Network Security Administration. Other criteria will apply in areas such as regional accreditation, outreach, and student and faculty research.

The new CAE IA/CD criteria became available in June 2013 for CAEs designated as research schools (CAE-R) to begin the revised process. The rest of the current CAE institutions and any new candidates are using a new web application that became available in October 2013. Schools with current CAE-IA designations are in the midst of transitioning to the new criteria. The goal is to insure that the students graduating from CAE schools will have the best skills conceivable to meet the needs of the nation.

Resiliency

In addition to working toward the next generation of cyber-security standards and practices and educating the future cyber workforce, the NSA must try to achieve an acceptable level of resiliency for its systems and networks to ensure data and operations are protected from unforeseen attacks. Unfortunately, no matter how hard the agency works, critical government networks will likely continue to be contested. For that reason, design efforts must include a focus on recovery and reconstitution. Now is the time to build resiliency into the nation's systems because the network ecosystem is at risk. National security solutions and practices must be embedded in the enhanced systems…not as afterthoughts. Working toward an acceptable level of resiliency will limit the damage that could be inflicted by threat actors.

The fact is that even when the right level of confidence in cyberspace is achieved, threat actors will still persist, and even the best defenses will not always be perfect and prevent 100% of the attacks. Thus, resiliency is necessary to limit the damage of those attacks that get through, speed reconstitution, and insure that business functions can continue even in light of threats.

Conclusion

In summary, achieving an acceptable level of confidence in cyberspace requires the composition of solutions to attain desired security and usability and the building of those solutions into the nation's systems. Moreover, these systems must be designed so that attackers must compromise or evade multiple independent mechanisms to achieve their goals. NSA should collaborate with key industry entities to include Internet service providers, as well as hardware and software companies, to achieve a scalable, broad impact. These efforts must keep pace with the commercial market place.

In addition, the NSA must build close partnerships across educational entities, government, and industry to facilitate real-time sharing of actionable information. Initial efforts must be on finding automated solutions for easy problems so that subsequent efforts can focus limited resources on the hard ones. Finally, the NSA needs to plan for an attack by building resiliency into networks, never forgetting that an ever-determined adversary will evolve and change even as we build a defensible environment.

NSA will continue to use everything in its arsenal to protect and defend National Security Systems, but future success will be a direct result of partnerships with industry, academia, and government. The result of these partnerships will be a more secure cyberspace ecosystem. By continuing to focus on these areas, NSA is certain that achieving confidence in cyberspace, despite the complexities created by an ever-changing ecosystem, is a goal that can be realized.

References

Danyliw, R, Meijer,J & Demchenk,Y 2007, *IODEF, Incident Object Description Exchange Format*, December 2007, viewed 18 Feb 2014, <http:// tool.letf.org/html/rfc5070>.

Davidson, M & Schmidt, C 2014, *TAXII –Overview*, Version 1.1, .pdf, 13 January 2014, Mitre, viewed 20 February 2014 ,<http://taxii.mitre.org/specifications/version 1.1/#spec>.

Digby 2014, *Mobile industry statistics*, 2014, Digby, viewed 22 February 2014, <http://digby.com/mobile-statistics>.

Evans, D 2011, *The Internet of things: how the next evolution of the Internet is changing everything*, white paper, April, Cisco Internet Business Solutions Group (IBSG), viewed 14 March 2014, <http://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf>.

Gantz, J & Reinsel, D 2012, *The digital universe in 2020: big data, bigger digital shadows, and biggest growth in the Far East*, viewed 18 March 2014, <http://idcdocserv.com/1414>.

Mahdi, W 2012, 'Saudi Arabia says Aramco cyberattack came from foreign states', Bloomberg L.P, viewed 17 March 2014, <http://www.bloomberg.com/news/2012-12-09/saudi-arabia-says-aramco-cyberattack-came-from-foreign-states.html>.

NIST 2014, *SCAP specifications: SP800-123 rev.2, version 1.2* , Security Content Automation Protocol, September 2011, NIST Information Technology Laboratory, viewed 20 March 2014 , <http:// scap.nist.gov/revision/index.html>.

STIX Structured Threat Information eXpression 2014, Mitre, 11 March, viewed 21 February 2014, <http://stix.mitre.org.language/version 1.1>.

Worldometers: Real Time World Statistics 2014,*Computers sold this year worldwide*, viewed 22 February 2014, <http://www.worldometers.info.computer>.

Information Assurance Standards: A Cornerstone for Cyber Defense

M Boyle

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *NSA has a rich history of contributing to standards that enable cyber defense. This paper examines that history, tracing the evolution of NSA's involvement in the development of early commercial encryption standards, through its more public contributions to network security protocols, to its current efforts to promote and create cyber standards that support the Department of Defense's use of commercial products to protect classified information.*

Keywords*:* *Standards, Cryptography, Network Security, Commercial Security*

Introduction

The primary mission of the National Security Agency (NSA) Information Assurance Directorate (IAD) is protection of National Security Systems (NSS) and the information they contain. When NSA was established in 1952, it had the dual mission of performing Signals Intelligence (SIGINT) and Communications Security (COMSEC). While the focus of this article is how IAD work in standards supports Cyber Defense, a brief look at how IAD's mission has evolved is in order.

The initial IAD COMSEC mission was targeted to protect against an adversary that would simply collect communications and read them in the clear. There were no complicated trust relationships; rather, communications were point-to-point. Encryption would be between two entities or sites and would be performed by boxes whose only job was to encrypt all the traffic passing in one direction and decrypt all the traffic in the other. The only other security issues were physical and procedural (personnel). Needless to say, with a relatively small market for secure electronic communications, these boxes were generally built by the Department of Defense (DoD) or its contractors (who received stringent requirements and guidance regarding the final product).

In contrast, today's environment (for the U.S. Government (USG) and most other large enterprises) is orders of magnitude more complex. Systems are networked, with routers choosing the best paths over which to forward packets. In general, there is no simple 'red side' (plaintext) and 'black side' (cipher). In fact, security associations may be transitory, lasting just for a session or just for certain types of traffic. They may also be set up between entities that have not previously had contact or directly shared information with each other. Moreover, encryption can happen at many different layers of the protocol stack: between adjacent devices (link layer), between routers that are many hops apart (network layer), and between clients and servers (application layer). Many of the devices that enterprise traffic may traverse are not even

ISSN 1445-3312 print/ISSN 1445-3347 online

nominally owned by that enterprise. And, with the explosion in the Information Technology (IT) market generally and IT security in particular, the devices that perform the encryption are not usually produced by the government, but rather by commercial vendors who produce equipment and software for government (U.S., allies, adversaries) and/or commercial enterprises, and consumers alike.

The evolution of the IT environment—from simple, point-to-point encryption performed by special purpose equipment built by contractors to stringent government specification, to a layered topology of encrypted communications with more complicated decisions to be made about encryption and trustworthiness performed by commercial products built by vendors with a broader customer set that dwarfs any government—left the USG with a hard choice. It could either attempt to build or contract all the IT security capability that it required and try to embed it in an environment that was rapidly evolving (and continues to do so) or it could give up considerable control over the equipment that it would use to secure its networks and instead rely upon commercial products. To a great extent, the USG has done the latter. In support of DoD's efforts to use commercial IT products, IAD has responded by engaging industry through programs such as Commercial Solutions for Classified (CSfC) (for more about this program, see 'Outmaneuvering cyber adversaries using commercial technologies' later in this journal) and has also stepped up its participation in public standards-development organizations.

Among other things, this decision leaves the USG in a bind. It has clear requirements (both operational and security) for performing its mission. These requirements are not always the same as those of commercial enterprises and consumers (whether they should be, and if so, who is at fault are matters of unending debate). How then is the USG going to find products that will meet its requirements? Put another way, if a provider of IT equipment cannot be swayed by a USG contract to make substantial changes to its products (generally, they cannot), how can the department express its requirements in a form that might create change? One answer is to become involved in the creation of standards to which products are built.

Standards are an answer to the one problem, but they also address another. Like any consumer, the USG's customers want to be able to select among a variety of vendors and products. Lack of standardization makes this difficult. This analogy illustrates the problem. If every electric company had its own interface and power characteristics, consumers who wished to populate their homes with refrigerators, ovens, washing machines, etc., would be locked in to one (or a small number) of vendors that built products conformant with each consumer's electric company. This would not be a satisfying situation for the consumer, who would likely suffer both in the price of appliances and in the features provided. So standards provide a mechanism for specifying a set of requirements, while at the same time allowing the customer the advantages that come from interoperability, namely choice and competition.

A Brief Aside on Policy

To explain IAD's role in setting requirements for the security of sensitive government networks, it must be noted that National Security Directive 42 designates NSA as the National Manager for National Security Systems (NSS) (Office of the President 1990). NSS are information systems, such as those used for intelligence activities or command and control of military forces, which require special protections (Barker 2003). NSA's role is to prescribe the appropriate protections

for NSS. In support of that role, IAD works to ensure that products (either commercially produced or created specifically for the government by its contractors) are available to provide that protection.

In addition to its NSS responsibilities, NSA has an advisory role with the National Institute of Standards and Technology (NIST). NIST is responsible for setting standards for the rest of the federal government information systems (that is, the non-NSS ones). *The Computer Security Act of 1987* (which was superseded by the *Federal Information Security Management Act of 2002*) describes NIST's role for setting these standards but also states that NIST should be "drawing on the technical advice and assistance (including work products) of the National Security Agency, where appropriate" (Committee on Science, Space, and Technology 1987). This advisory role was formalized in a Memorandum of Understanding between NIST and NSA in 1989.

Given the foregoing, how has NSA engaged industry through standards-development organizations, and what has the effect of that engagement been? The answer is that NSA's involvement has evolved over time, both in method and also with regard to the type of standard that it has prioritized.

A Suite of Algorithms

A good starting point is the development and adoption of the Data Encryption Standard (DES). (The Data Encryption Standard has been retired. See United States Department of Commerce 1999 for the last version of *Federal Information Processing Standard (FIPS) 46*.) The move to create a public standard for encryption by the then National Bureau of Standards (predecessor to NIST) was a watershed moment in public cryptography. For the first time, there was an algorithm that was publically documented and intended for use by a community broader than government (initially, the banking community). First proposed by researchers at IBM, DES was adopted by NIST in 1976 (as FIPS 46) and went on to become widely used. Prior to this, NSA had examined DES and found that a critical component of the algorithm, a set of tables for performing substitutions (the S-Boxes), were not well chosen (George 2011). NSA took the bold step (for that time) of recommending that the S-Boxes be changed and even provided new ones, which were eventually used. Naturally, this change was heavily scrutinized by the public. NSA did not take credit for the change, mostly due to the ethos at the time that NSA stood for No Such Agency. The changes led the nascent community of public cryptographers to develop new techniques that shed light on the selection criteria for the new S-Boxes.

Over time cryptographic algorithms fell to advances in cryptanalysis and inevitably to Moore's Law. Eventually, DES faced the end of its life as a viable encryption algorithm. NIST adopted the Advanced Encryption Standard (AES) after a lengthy competition, in which cryptographers from around the world submitted candidate algorithms. NSA did not submit an algorithm. It did provide extensive technical support to NIST by examining the algorithms and offering technical guidance.

In the early 1990's, NIST decided to standardize a hash algorithm, which, among other things, would be a necessary component to creating a digital signature algorithm. NSA offered up an algorithm, initially called the Secure Hash Algorithm (SHA). This algorithm was well on its way to becoming a standard when IAD cryptanalysts discovered a weakness. NSA quickly designed a

tweaked version of SHA that addressed the flaw. The original version was renamed SHA-0, and the newer version, SHA-1 became the U.S. national standard in 1995. (See National Institute of Standards and Technology 2014, *Federal Information Processing Standards FIPS publications* for the current versions of standards referenced in this section.)

In the early 2000's, NIST saw the need to create new hashes that could support higher security levels. NSA had ongoing research in this area and offered a suite of algorithms that produced hash values of lengths (256, 384, and 512 bits) greater than SHA-1 (allowing for the possibility that, if designed correctly, they could offer greater security). This family of algorithms was adopted as SHA-2. To date, there have been no successful attacks against the SHA-2 family.

NSA designed both the original version of the Digital Signature Algorithm (DSA) and the Elliptic Curve DSA, which were adopted as NIST standards (FIPS 186). NIST has also standardized other public key systems, including RSA, which are in common use today.

It is worth noting that in the mid-2000's NIST decided that it would be prudent to establish another suite of hash algorithms. At the time, there was a frenzy of public cryptanalysis of SHA-1, which was yielding results. The feeling was the SHA-2 family was similar enough to SHA-1 that the work being done might be adapted to SHA-2. This has not turned out to be the case (as always, the caveat must be 'to date'). At any rate, NIST held a public competition to create another family of hashes. Submissions poured in from around the world (the deadline was 31 October 2008), and the winning algorithm, Keccak, also referred to as SHA-3, was selected in October 2012. NSA did not submit an algorithm, instead serving as a technical evaluation resource, performing cryptanalysis on the submissions and offering input to the NIST process.

Having provided technical support to the development and selection of a set of algorithms suitable for protecting government communications (which are now used to protect corporate and private transactions as well), NSA proceeded to adopt a subset of those algorithms that could be used in concert to protect National Security Systems (Committee on National Security Systems 2012). In time, this set of algorithms was referred to as Suite B. It consisted of AES, SHA-2, Elliptic Curve Diffie Hellman, and Elliptic Curve DSA (each of two different security strengths). This was a necessary, but not completely sufficient step in securing communications for NSS. The next aspect of the problem to tackle was to select a set of protocols that the suite could be embedded in and to work out the details of the key management.

Securing Communications on the Network

With a suitable set of cryptographic algorithms either in hand or under development, NSA considered the landscape in which their customers were beginning to operate. In the 1990's, enterprises started to become reliant on email to conduct day-to-day business. Also, in parallel to the explosion of the World Wide Web, business processes began to acquire web interfaces. Combined with the move of communications to the Internet, it was abundantly clear that a set of security protocols was needed to leverage the cryptographic suite. Luckily, the private sector was well on its way toward solving this problem after early investment from the DoD in developing the underlying network (specifically the Advanced Research Projects Agency (ARPA), which was renamed the Defense Advanced Research Agency (DARPA) in 1996). Early versions of

Internet Protocol Security (IPsec) and Transport Layer Security (TLS) existed. Somewhat fortuitously, NSA had contributed to the development of IPsec.

In the late 1990's, employees in NSA's research organization began to consider the problem of establishing secure tunnels on a packet switched network. They foresaw the need for a protocol that would not only encrypt packets between two endpoints, but would also allow for flexible authentication and for creating key material on the fly. The result, Internet Security Association and Key Management Protocol, ISAKMP, was a protocol that provided authentication as well as negotiation of algorithms and derivation of keys that could then be used in a subsequent encryption protocol (Maughn *et al*. 1998). ISAKMP was submitted to the Internet Protocol Security (IPsec) Working Group at the Internet Engineering Task Force (IETF) and influenced the eventual design of the Internet Key Exchange (IKE) (Harkins & Carrel 1998) and IPsec (Kent & Atkinson 1998). In particular, the two-phase negotiation in IKE is based on the ISAKMP design.

The IETF had IPsec. They also had developed a protocol for encrypting traffic between a client and server, Transport Layer Security (TLS) (Salter & Housley 2012), which was based on the earlier Secure Sockets Layer (SSL) protocol designed by Netscape. TLS has been widely deployed as the primary protocol for authenticating web servers and protecting the traffic between the server and browsers. It is likely the most widely deployed network security protocol (and easily the most recognizable to the casual computer user).

The IETF also had a security format for email messages: Secure Multipurpose Internet Mail Extensions (SMIME) (Ramsdell 1999). SMIME allows a user to encrypt and sign his email with common commercial algorithms. It has been implemented widely, especially in email tools intended for corporate environments. The initial digital signatures (for signing the email message) and key management were based on the RSA public key scheme.

NSA recognized the utility of these protocols. But there were some problems, principally concerning the cryptographic strength of the algorithms specified. In particular, the public key sizes that were used would not be suitable for protecting National Security Information (NSA deemed 128 bits of security to be sufficient strength for protecting information up to the SECRET level; typical implementations of RSA in commercial protocols offered much less).

NSA's approach to the problem was two-fold. First, it was necessary for two parties involved in communication to be able to negotiate an acceptable set of cryptographic algorithms. All the protocols mentioned above had that capability to some extent. However, they generally were not very extensible (that is, it was not possible to add new features, including new algorithms), and they often had some algorithms hard-coded in certain parts of the protocol (SSL, the predecessor of TLS, did not allow negotiation of hash algorithms; neither did TLS 1.1). NSA had a preferred suite of algorithms, NIST-approved and ready to go. It needed to be able to negotiate them in each protocol of interest, including having assigned identifiers that could be used in the negotiation (this is referred to as crypto-agility). Further, NSA needed to publish reference documents that would tell implementers what set of algorithms should be used and how to use them together in each protocol.

The participants in the IETF were in agreement that protocols should allow new algorithms to be specified and negotiated. Besides satisfying the desire to support algorithms for various communities (including other national algorithms), adding this type of flexibility also enabled a protocol to negotiate new algorithms if problems were found with the existing ones (without having to rewrite the standard). No engineers want to intentionally paint themselves into a corner. Over time, IPsec, TLS, and SMIME all became highly crypto-agile. While this was happening (with NSA as one of many supporters), NSA began to author Internet Drafts in the IETF to support Suite-B implementations. In some cases, only an Internet Draft with the appropriate algorithm identifiers was produced. In other cases, another draft that described particular requirements for a Suite-B implementation was also created. With multiple drafts available (Salter & Housley 2012; Burgin & Peck 2011), vendors had the information they needed to include Suite-B support in their products.

NSA's efforts helped raise the cryptographic strength of these Internet protocols. AES implementations replaced DES or Triple DES. SHA-2 replaced SHA-1 in some places where digital signatures were required (SHA-1 has lingered in some areas, particularly in the public key infrastructure that all three protocols depend upon). The one area where Suite B has lagged has been in replacing RSA and Diffie Hellman with Elliptic Curve Cryptography. There are many reasons, but patents related to the technology are probably the main one. This patent situation has led to the continued use of RSA at security strengths that are less than that of the rest of the algorithms in the implementation (due to the inefficiency of RSA at the required strength).

Improving the cryptography in network protocols addressed the problem of protecting customer data. But as networks begin to offer richer interfaces, the nodes on the networks become more vulnerable. The next challenge was to address this new threat.

Hardening the Platform, Making the Network Manageable

As terminals connected to mainframes gave way to workstations with greatly increased functionality, security of those workstations became a serious concern. The history of attacks on personal computers is well documented. Perhaps not as well examined is the movement of those attacks, from applications, down through the Operating System, and eventually to the BIOS. In recent years, there has been a great deal of effort from industry to combat this problem. These efforts have focused on 1) making the BIOS harder to corrupt, 2) detecting changes to the BIOS, and 3) enabling recovery from a corrupted BIOS.

IAD became concerned in the early 2000's about the possibility of an attacker making unauthorized changes to the BIOS of a personal computer and leveraging that to have unrestricted (and almost undetectable) access to the information on that machine. Later, that concern expanded to the notion that an adversary could simply corrupt the BIOS on a machine (or even a whole network of machines) and cause a denial of service from which it would be very difficult to recover. With these concerns in mind, and with some internal research under its belt, IAD turned to the Trusted Computing Group (TCG) as a venue to work on this problem.

The TCG (a consortium of leading IT companies) had been developing specifications for a Trusted Platform Module (TPM) since the mid 1990's. At the time NSA joined TCG, hardware compliant with the TPM 1.2 specification was already widely deployed on personal computers

around the world. It was not being widely used. IAD's first concern was to make sure that the next-generation TPM would meet all USG requirements, to include cryptographic algorithms that would be acceptable on USG networks, as well as robust mechanisms to counter the threats already mentioned and to enable emerging use cases.

As IAD's involvement in the TCG deepened, it became apparent that there were plenty of well-qualified people in the TCG to deal with technical issues, but there was a real need for more people who could speak to how they would use Trusted Computing products if they existed, and to advocate for sufficient security robustness (not that the NSA participants at TCG ignored the technical debates; they could not help being involved there as well). One message that NSA regularly articulated was that it was not enough to have products that implemented TCG standards. What was needed was an ecosystem of products that could solve customer problems. More attention needed to be paid to two areas

- stitching various specifications together (for instance, so that an encrypting hard drive could use the TPM for key storage or a network decision point that had evaluated a machine joining the network could share information about that machine with other security decision makers on the network) and
- testing interoperability of the various components of a solution (to ensure that when products that implemented TCG standards were purchased, they could be readily configured to work together in common scenarios).

IAD recognized early on that one of the big challenges in network security is that network administrators really have too many tasks to do just keeping a network running, and too much data to analyze to have any hope of responding to security threats. As a rule, IAD's advice is that a well-run, promptly patched network is more resilient to attack (see, for example, NSA's *Manageable network plan*). Keeping all the systems on a network up to date and operating within policy is a daunting task, given operational demands. If only there were a way to take some of the mundane tasks out of the administrators' hands, thus freeing them to perform tasks that only a well-trained human could execute.

Along with the Department of Homeland Security, NIST, and partners in industry, NSA created a set of specifications for communicating information about the configurations of machines and vulnerabilities that have been discovered. The intent was to allow machines to have conversations about their configuration, their patch status, and emerging threats or vulnerabilities to look for. Providing this capability would not only clear out the underbrush for the network administrator, but it would also allow the network to respond swiftly to new vulnerabilities, by comparing the necessary conditions for exploiting the problem to the actual configuration of the machines on the network. The result of the initial effort was the Security Content Automation Protocol (SCAP), which was a set of related schema for communicating configuration, vulnerabilities, and the like. (For more information about SCAP, see scap.nist.gov.) SCAP represented a first step in dealing with the problem, seeing some adoption, and spawning an effort in the IETF to create a more comprehensive solution that would include not only the schema, but also the networking standards to carry the information.

Note on Making Standards Have an Impact

When NSA participates in a commercial standards organization, it can play the role of a subject matter expert on a wide range of topics (its core competence, cryptography, but also maintaining network security and operating a large scale network in general). In contrast to its participation in international or USG-sponsored bodies, NSA also fills a niche that is often not well-occupied in these venues, namely as a customer voice that can directly state requirements for a decent-sized sector of the market. While the typical attendee at a commercial standards meeting (for instance, the IETF or TCG) is either directly or indirectly representing the interests of a vendor who will build products that meet the resulting standard, NSA is able to present use cases and make requests for levels of security, resilience to attack, and other features.

Complementary to this role, NSA has an interest in making sure that products actually show up to be purchased (standards are important; products implementing those standards correctly are essential). In the past, ensuring the availability of products would have been handled by contracting for equipment to be specially built for the USG that would implement all stipulated requirements. In the case of commercial products, the influence is much more indirect, based on sending clear signals to industry that if a product is built that meets a given set of requirements, including implementation of the correct standards, purchases will be made by the DoD and other parts of the USG. If partners in the governments of other nations can agree to this, the case for implementing a standard is even better. Thus NSA tries to do the following when a standard has been created that is deemed to address its requirements.

1) Work with partners in the USG and with foreign government partners to achieve consensus on implementing the standard.
2) Cite the standard in procurement language, where appropriate, so that purchasers are encouraged to buy products that meet the standard.
3) Create requirements in appropriate Protection Profiles that can be met by correctly implementing the standard.

It might not be obvious, but getting equipment that meets a certain standard into the field is not the end of the story. Standards are not always implemented correctly, which can lead to poor security or interoperability issues (in the case where two vendors do not agree on what the standard means). IAD tries to address this through its participation in Common Criteria and its reliance on evaluations against suitable Protection Profiles for vetting products used in its CSfC initiative (NSA/CSS 2014). As well, IAD engages actively with NIST to ensure that FIPS 140 and the Crypto Module Validation Program (CMVP) (NIST 2014, *Cryptographic Module Validation Program (CMVP)*) offer a sound avenue for evaluating the cryptographic pieces of CSfC components. NSA has worked with NIST to create testing language for its Protection Profiles that is consistent with CMVP test requirements.

It is often the case that features in a product that are based on a standard (or more generally that provide a given capability) do not actually get turned on and used. This has certainly been the case for TPMs, until recently. Hundreds of millions of computers have been shipped with TPMs on the mother board. However, most are not activated, let alone used in any way. NSA has encouraged the use of TPMs through procurement language. Recently, IAD issued an IA Advisory to its DoD customers, describing how TPMs can be used to enhance the security of certain applications, and encouraging that they not only be purchased, but also actually turned on

and leveraged for these applications. This use will be contingent on a CMVP validation as well as Common Criteria evaluation against a TPM Protection Profile, which is currently being vigorously developed.

Future Work

Working through standards-development organizations is a long-range activity as a given standard often takes three (or more) years to be proposed, argued, agreed upon, and published. While it is difficult to say which technical problems will be most pressing ten years from now, some trends are evident. Enterprises (as well as consumers) are likely to be more concerned with small, mobile devices than with large boxes sitting under desks. There will be vastly more data generated by familiar objects that are not currently participating on the network. There will, thus, be an increase in data moving around the network, and generally that data may be stored outside the enterprise (in fact, a lot of data processing may occur on machines not directly controlled by the enterprise). And, while it is too late to call online commerce a future trend, recent, well publicized problems with consumer data are likely to continue. In light of all this, here are some thoughts on how standards could be developed to cope with the evolving environment.

Industry is still trying to decide how to secure mobile devices (for instance, phones and tablets). The hardware-rooted trust that has been worked out for PCs and servers is still being debated for mobile platforms. IAD has asked, via its Protection Profile for Mobile Device Fundamentals (NIAP 2014), for industry to provide hardware protections for keys on mobile devices, and would likely ask for more hardware protections if industry consensus is achieved. Work is currently being done in the TCG and in the Global Platform Alliance to specify this.

As every conceivable device hooks up to the Internet (already DVRs, soon refrigerators), there is a need to secure communications with those appliances. It may be the case that authentication and authorization will be more important than encryption (while individuals may not want others seeing how much ice cream they are consuming, they certainly do not want others to have the ability to take ice cream off their grocery lists or turn the temperature up in their refrigerators). Standardizing on a set of flexible authentication and authorization mechanisms will be important.

Authentication of people on the Internet (so they can make purchases or access services) also needs to be worked out. To be clear, there are plenty of mechanisms out there. Settling on a scheme that does not involve the consumer providing a credit card number to every online merchant is the challenge. There are already solutions involving third-parties with which the user already has a relationship. A standards-based solution that allows more flexibility would be useful.

As data storage and processing is out-sourced more and more, the need to manage the security of that data will increase. Currently, standards are lacking for labeling data as well as ensuring that the proper security mechanisms are being applied. At some point, a framework for negotiation of levels of service will likely emerge, but many enterprises will need assurance that the data is being hosted according to their security requirements. Defining a protocol that provides those assurances will be challenging.

Conclusion

NSA's involvement in standards development has evolved over time, both in the nature of its participation (first as an indirect, largely silent partner, focused only on technical issues to a direct participant, speaking as a customer voice) and its subject-area focus (cryptography, secure network protocols, more general network security and manageability). NSA's continued participation will be focused on the emerging environment, with an eye toward stating requirements and proposing solutions that benefit not only the USG, but also the nation's critical infrastructure and online commerce.

References

Barker, W 2003, *Guidelines for identification of Information Systems as National Security Systems*, viewed 25 March 2014, <http://csrc.nist.gov/publications/nistpubs/800-59/SP80059.pdf.>.

Burgin, K & Peck, M 2011, *Suite B profile for Internet Protocol security (IPsec).* viewed 25 March 2014, <http://tools.ietf.org/html/rfc6380>.

Committee on National Security Systems 2012, *CNSS Policy No. 15: National Information Assurance Policy on the use of public standards for the secure sharing of information among National Security Systems*, viewed 25 March 2014, <https://www.cnss.gov/CNSS/issuances/Policies.cfm>.

Committee on Science, Space, and Technology 1987, *Computer Security Act of 1987*, report, Committee on Science, Space, and Technology, United States House of Representatives, <http://csrc.nist.gov/groups/SMA/ispab/documents/csa_87.txt>.

George, R 2011, 'NSA's role in the development of DES', speech, *Proceedings of RSA conference, topics in cryptology – CT-RSA 2011*, A Kiyias (ed.), LNCS 6558, San Francisco, California, United States, p. 120.

Harkins, D & Carrel, D 1998, *The Internet Key Exchange (IKE)*, viewed 25 March 2014, <https://www.ietf.org/rfc/rfc2409.txt>.

Kent, S & Atkinson, R 1998, *Security architecture for the Internet protocol*, viewed 25 March 2014, <http://tools.ietf.org/html/rfc2401>.

Maughn, D, Schertler, M, Schneider, M, & Turner, J 1998, *Internet Security Association and Key Management Protocol (ISAKMP)*, viewed 25 March 2014, <http://www.ietf.org/rfc/rfc2408.txt>.

National Information Assurance Partnership 2014, *NIAP approved protection profiles*, 20 March, National Information Assurance Partnership: Common Criteria Evaluation & Validation Scheme, viewed 25 March 2014, < https://www.niap-ccevs.org/pp/>.

National Institute of Standards and Technology 2014, *Cryptographic Module Validation Program (CMVP)*, Computer Security Division, Computer Security Resource Center, 26

February, NIST Information Technology Laboratory, viewed 25 March 2014, <http://csrc.nist.gov/groups/STM/cmvp>.

—2014, *Federal Information Processing Standards FIPS publications*, Computer Security Division, Computer Security Resource Center, 18 March, NIST Information Technology Laboratory, viewed 25 March 2014, <http://csrc.nist.gov/publications/PubsFIPS.html>.

National Security Agency, *Manageable network plan*, technical report, NSA, viewed 25 March 2014, <http://www.nsa.gov/ia/_files/vtechrep/ManageableNetworkPlan.pdf>.

—/Central Security Service 2014, *Commercial Solutions for Classified Program*, 20 February, National Security Agency/Central Security Service, viewed 25 March 2014, <http://www.nsa.gov/ia/programs/csfc_program>.

Office of the President of the United States 1990, *National Security Directive 42*, viewed 25 March 2015, <http://www.fas.org/irp/offdocs/nsd/nsd42.pdf>.

Ramsdell, B. 1999, *S/MIME Version 3 message specification*, viewed 27 March 2014, <http://tools.ietf.org/html/rfc2633>.

Salter, M & Housley, R 2012, *Suite B profile for Transport Layer Security (TLS)*, viewed 25 March 2014, <https://tools.ietf.org/html/rfc6460>.

United States Congress 2002, *Federal Information Security Management Act of 2002*, pdf, viewed 25 March 2014, <http://csrc.nist.gov/drivers/documents/FISMA-final.pdf>.

United States Department of Commerce 1999, *Data Encryption Standard (DES)*, archived publication, 25 October, Federal Information Processing Standards Publication, National Institute of Standards and Technology Information Technology Laboratory, <http://csrc.nist.gov/publications/fips/fips46-3.pdf>.

Cyber-Mugging: Summary and Analysis of a Simulated ICS/SCADA Attack

PJ DeSantis

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *In a representative Industrial Control System (ICS)/Supervisory Control And Data Acquisition (SCADA) laboratory environment, a simulated cyber attack suggests that an attacker with a low to moderate level of technical proficiency may utilize common, publicly-available tools and techniques to obtain complete control of the ICS environment. The cyber-physical relationship between information systems and industrial machinery has created environments where limited resources may be leveraged to trigger significant physical effects. The feasibility that such an incident has the potential to cause significant disruptive effects directly challenges the current paradigm that state-level resources are required to inflict catastrophic results.*

Keywords*: Industrial Control Systems, ICS, SCADA, Cyber Security*

Introduction

Industrial Control Systems security

The term Industrial Control System (ICS) describes information systems that are used to control industrial processes. A subset of ICS, Supervisory Control and Data Acquisition (SCADA) systems employ centralized data collection, monitoring, and control capabilities to manage smaller, geographically dispersed control systems. Collectively referred to as ICS/SCADA, these systems are essential to the manufacturing, production, and distribution processes of nearly all industries, especially in critical infrastructures such as electric, water, and oil and natural gas (National Institute of Standards and Technology 2013; Stouffer,Falco & Scarfone 2011).

As information systems, ICS utilize many of the same hardware, software, and communications capabilities of traditional Information Technology (IT) systems. The distinguishing characteristic of ICS is the capability to directly affect the physical world. This ‘cyber-physical’ relationship facilitates an abundance of beneficial technological capabilities, including process automation, energy production and distribution, and rapid global transportation. However, the effects of the corresponding risks inherent to ICS range from inconvenient to environmentally devastating. In extreme cases, the health and safety of human lives may be endangered (Zhu, Joseph & Sastry 2011).

In pursuit of efficiency and profitability, ICS are often designed and implemented with availability of the system as the top priority (Weiss 2010, pp. 208-210). This demand for availability drives unique performance and reliability requirements, frequently resulting in the use of operating systems, applications, and devices that may be considered exotic to many

traditional IT administrators. In the classic battle between usability and security, availability requirements in ICS will always tip the scales in favor of usability, leading to the ongoing deployment of systems which have the capability to exert control over physical machinery, yet were designed with little to no cyber-security considerations.

The potential for cyber attacks of ICS to result in physical damage is not a new concept and has been extensively demonstrated and documented. As compiled and outlined in the *National Institute of Standards and Technology (NIST) Special Publication 800-82, Guide to Industrial Control Systems Security* (Stouffer, Falco & Scarfone 2011), there have been a number of reported incidents during which attackers intentionally targeted ICS, malware unintentionally caused collateral damage to ICS, and unintentional internal security consequences resulted from seemingly benign activities such as security testing or configuration changes. Referenced by *NIST SP 800-82*, Sandia National Laboratories (Duggan, *et al*. 2005) describes incidents in which a simple ping, a standard network diagnostic utility, caused a robotic arm to activate from standby mode in one instance, and shut down a manufacturing system in another.

Although there have been numerous reports of cyber events targeting ICS, as well as detailed and elaborate analysis of some of the more high-profile events (Weiss 2010, pp. 107-122), and an abundant number of reported 'close calls' that provide a reminder of the risks posed by ICS control over physical components, there has yet to be an incident which resulted in catastrophic failure or loss of life that was definitively proven to be the result of a cyber attack.

Objectives of the cyber-mugging exercise

It is safe to assert that the absence of significant destructive cyber attacks targeting ICS is in no way a result of common ICS security practices. However, there is a debate as to just how low the bar is with regard to attacker skill level and resource requirements needed to execute destructive cyber attacks against ICS. Inspired by a perceived disparity between the assertions of security specialists and researchers on one side and industry and government on the other, the cyber-mugging exercise was conducted to simulate an opportunistic, 'target-of-opportunity' attack in which an actor with little to no ICS knowledge gained access to an ICS environment and attempted to destroy ('mug') the system.

To conduct this exercise, an information security analyst was tasked with performing a system-level vulnerability assessment and proof-of-concept intrusion of an ICS by an attacker with no technical or functional ICS knowledge. The 'attacker' was presented with a representative ICS environment and given limited time, open source tools, public vulnerabilities and exploits, and widely-known techniques, with the primary objective of determining if, and to what extent, critical ICS functionality is available and subject to abuse by an attacker without the use of ICS-specific knowledge and exploits.

Limitations on scope and capabilities of the attacker

The simulated attack was not intended to demonstrate advanced ICS exploitation techniques, extensive post-exploitation activities, or covert persistence strategies. Discovery of new ('0-day') vulnerabilities and exploit development were considered beyond the scope of the exercise. The attacker had professional information-security vulnerability-assessment and penetration-testing experience, but lacked any ICS expertise or knowledge of the target system configuration.

Finally, the attacker was restricted from utilizing commercial tools and from directly connecting to the ICS network.

Technical Environment

The test environment was procured from an authorized Siemens distributor and system integrator in 2012 to support vulnerability analysis of ICS components, software, and networks. This system was comprised of TCP/IP Ethernet networks and PROFIBUS (a control systems protocol) sub-networks that connected ICS devices, communications equipment, and Microsoft Windows XP-based servers and workstations loaded with Siemens PCS7 (WinCC and STEP 7) software. ICS devices connected to the network included one Siemens S7-317 and two S7-414 Programmable Logic Controllers (PLCs) with corresponding Ethernet and PROFIBUS communications modules, multiple Siemens SCALANCE X-200/300-series Ethernet switches, digital and analogue input-output (I/O) modules, and an ICS process simulator.

The network topology was logically segmented and employed a Siemens SCALANCE S-612 firewall to restrict communication between the 'external' and 'internal' (ICS) networks. All ICS devices were on the internal network. The external network contained only a single client workstation, the external firewall interface, and the attacker's workstation.

To best represent a realistic ICS environment, the system configuration was intentionally maintained in an 'as-delivered' state. As a result, the target system lacked implementation of information-security 'best practices' such as password-complexity requirements, network-intrusion detection systems, physical network segmentation, or up-to-date system patching and configuration management. While the security configuration was admittedly poor, it was designed and configured by manufacturer-authorized system-integration professionals and emulated a realistic ICS environment.

The industrial process being simulated by the ICS was modeled after a subset of equipment, materials, and processes involved in the production of ammonium nitrate, a compound used in products including agricultural fertilizers as well as explosives. This simulation included multiple raw material inputs, process reactors, and product storage facilities. The process control objects and sensors, used to measure and report values for pressure, flow, and levels, were simulated utilizing the Siemens SIMIT virtual commissioning simulation framework. Visualization of the process, including alarms, interlocks, and automated actions, was provided via the Siemens PCS7 WinCC Human-Machine Interface (HMI) package. The control system was configured and programmed utilizing the PCS7 Engineering package.

The simulated attack was conducted in January of 2013 by a single attacker and was witnessed by a control systems engineer and an IT administrator.

Attack Narrative

Tools and techniques of the attacker

All of the software tools utilized in this exercise are free, open source, and well documented. It should be highlighted that, although free to obtain and use, they are nonetheless professional-quality products. As with many information-security applications, these tools are developed and

distributed for legitimate personal or authorized professional use. However, there are no functional restrictions preventing malicious use.

The entirety of the attack was conducted using tools that were included in Back-Track (recently rebranded as Kali) Linux, distributed by Offensive Security. This includes Rapid7's Metasploit Framework (www.metasploit.com), Nmap (nmap.org), and a handful of Linux command-line utilities. Dedicated vulnerability scanners were not utilized in the simulated attack.

The attack was conducted using simple and common methods of information gathering, vulnerability discovery, exploitation, and post-exploitation activities. An attacker with a low to medium level of technical expertise would be able to replicate the techniques used in the attack without difficulty. Mastery of these techniques would be considered 'entry-level' for a professional information-security assessor.

Reconnaissance, scanning, and service enumeration

From a dedicated attack workstation on the network 'outside' of the firewall, the attacker initiated reconnaissance of the environment by launching a quick Nmap scan to discover hosts sharing the external network. Nmap utilizes a variety of built-in scripts, known as the Nmap Scripting Engine (NSE), to interrogate and enumerate the details of running services. These scripts may provide the attacker with configuration details or even positive identification of vulnerable services. The initial scan results indicated that there were two hosts with which the attacker could communicate: the firewall and a single workstation.

The attacker chose to focus on the workstation due to the likelihood that it was a softer target than the security-focused firewall device. With only a single target, and no suspected intrusion-detection capabilities in the external network, the attacker unleashed a full Nmap scan against the workstation, probing all 65,535 ports for each of the TCP and UDP protocols, and running enumeration scripts on services discovered during the scan. The scan took less than 40 seconds to complete and returned 15 pages of new information related to the target workstation, which is significantly more than the few lines of results that may be returned when scanning a highly-secured host.

The large volume of information returned from the scan was primarily due to the Simple Network Management Protocol (SNMP) service configuration. This service was discovered as part of the UDP scan and was found to be configured with weak, although common, values for the public (read) and private (read/write) community strings. Community strings are SNMP's functional equivalent of authentication. The values for the community strings on the target workstation were 'public' and 'private' respectively, which is similar to using a password of 'password'.

Although SNMP disclosed a bulk of the scan's findings, only 6 seconds were required to fully interrogate the SNMP service. With full access to this service, the Nmap results included details of the Windows version and service pack level (Windows XP SP3), user accounts, all running network services with respective port numbers, all established network connections, details of physical and logical network interfaces, all Windows data shares, all running processes (name,

process ID, path), all installed software applications and version details, and all installed Windows patches and hotfixes.

The value of the information disclosed by SNMP cannot be overstated. The attacker was now aware of nearly every detail of the target workstation by leveraging functionality of a legitimate service; no exploit was needed. Of particular interest was the list of installed Windows patches, which disclosed to the attacker the specific Windows' vulnerabilities that had been remediated and, therefore, which exploits should work and which would fail. For example, although the target was a Windows XP SP3 system, SNMP disclosed that 'Security Update for Windows XP (KB958644)' was installed, indicating that a very reliable post-SP3 exploit would fail (Microsoft Corporation 2008).

When configured properly, SNMP does not necessarily expose a system to risk of exploitation. SNMP is often targeted in initial stages of an attack to obtain information that may disclose other system vulnerabilities. Older versions of SNMP (v1 and v2c) do not use encryption; and, even if configured with unique community strings, these strings will traverse the network as plaintext and may be captured by an attacker. Version 3 of SNMP has added encryption and authentication capabilities, but it is more complex to properly configure, and vendors have even shipped devices with flawed implementations, such as a recent vulnerability in Siemens' implementation of SNMPv3 that allows attackers to execute SNMP commands without credentials (ICS-CERT 2013). The prolonged device lifecycle in ICS dictates that even newer SNMPv3-enabled devices will continue to support v1 and v2c of the protocol to provide backwards compatibility.

Overall, the initial scan revealed more than two dozen network services, including Microsoft's SQL Server (MSSQL) database application. SNMP provided the attacker with the MSSQL version (2005 - 9.00.4053.00), database name (WINCC), and listening port (1031). Observing that MSSQL was listening on port 1031, which is anomalous from the standard port of 1433, the attacker suspected either configuration weaknesses or a deliberate attempt at hiding the service on a non-standard port. Using a Metasploit auxiliary module to enumerate further details of the MSSQL service, the attacker discovered that both remote access and xp_cmdshell were enabled. MSSQL instances that accept remote connections are a common target of exploitation due to the high likelihood that the database user is a Windows SYSTEM-level account. xp_cmdshell is a stored procedure that takes input to the database and passes it to the underlying operating system for execution (Kennedy *et al*. 2011, p. 79). Effectively, a remotely-accessible MSSQL instance running with SYSTEM-level privileges and xp_cmdshell enabled may provide an attacker with total control over both the database and the target operating system.

Exploitation

Gaining access to the external workstation

Identifying MSSQL as a potential foothold on the target system, while also inferring that the poor SNMP configuration could be indicative of pervasive weak credential use, the attacker initiated a brute-force attack on the MSSQL service. The multi-threaded, brute-force authentication tool Hydra (www.thc.org/thc-hydra) provided the attacker with the capability to attempt tens of millions of username/password combinations per day. In this exercise, it took only seconds for Hydra to report that the user Administrator (earlier disclosed by SNMP) had a

password of 'Administrator'. With no brute-force countermeasures in place, even a complex password would likely have been revealed in hours to days.

In possession of valid credentials, the attacker utilized Metasploit's mssql_payload module, which makes use of the xp_cmdshell stored procedure to inject a Meterpreter payload into memory. A powerful feature of the Metasploit Framework, Meterpreter is a payload that is executed by the victim when it is exploited (Kennedy *et al*. 2011, pp. 79-80). Running in memory and obfuscated by Metasploit's payload encoders, it often avoids being detected by anti-virus applications. Meterpreter provides the attacker with a console that enables automated execution of a wide variety of common post-exploitation activities. For example, extracting Windows NTLM password hashes from a victim is as easy as executing the 'hashdump' command at a Meterpreter prompt (Kennedy *et al*. 2011, p. 95). These password hashes may be cracked and/or used in 'pass-the-hash' attacks on other Windows systems throughout the network.

Upgrading access to include graphical capabilities

Legitimate users of ICS/SCADA systems, such as operators and engineers, rely heavily on the Human-Machine Interface (HMI) to interact with the system. The HMI allows users to monitor and configure set points, control algorithms, and adjust parameters in controllers, while also displaying process status information and historical information. The HMI is a point-and-click graphical tool, necessitating that an attacker obtain Graphical User Interface (GUI) functionality before accessing the HMI's capabilities.

Meterpreter is a command line utility and does not, by itself, provide GUI access. It is possible to inject and connect to a Virtual Network Computing (VNC) server via the existing Meterpreter session, which is helpful in situations where valid user credentials are unavailable or network security devices prevent inbound Remote Desktop Protocol (RDP) connection requests. However, VNC would be available for only as long as the Meterpreter session remained active. To provide persistent GUI access, the attacker used Metasploit post-exploitation modules (enable_rdp and incognito/add_localgroup_user) to enable the target's Remote Desktop service, open the appropriate port in the Windows firewall, and add the Administrator user to the 'Remote Desktop Users' group. With the target now configured to accept Remote Desktop connections, the attacker logged in to the victim's machine using the rdesktop (rdesktop.org) utility for Linux and gained access to the graphical HMI for the ICS.

With GUI access to the workstation in the external network, the attacker was able to use the HMI to perform the functions of an operator-level user. While operators do not have the same level of control of the ICS as an engineer or administrator, functionality was available that had the potential for misuse. A single click on a graphical switch was all that was needed to change the state of a reactor from 'on' to 'off'. While nothing destructive had occurred, the attacker was able to use a workstation on the external network to affect change to the process using nothing more than a mouse-click.

Bypassing the firewall to access the ICS

In addition to operator-level HMI access, the compromised workstation was also hosting the firewall configuration files and application. The Siemens SCALANCE S-612 firewall

documentation advises that it be configured from the external network (Siemens AG 2005), giving the attacker the ability to use this access to modify the firewall settings in a way that would allow unfettered access to the internal ICS network from the external environment. Using the Siemens Security Configuration Tool application that was installed on the compromised workstation, the attacker opened the firewall configuration file, which was named 'test'. When presented with a login challenge, the attacker entered the username 'test' and a password of 'test', which were valid credentials and allowed the attacker to load and modify the firewall configuration.

The intentional placement of the firewall configuration utility on a system external to the firewall exemplifies a vulnerability that the author speculates is the result of design choices made with the intention of keeping the management of what may be considered 'IT components' outside of the ICS. This may satisfy some system owners who wish to use the firewall as a demarcation point between what is considered the ICS and what is not. However, it also provided the attacker with the ability to modify firewall settings and gain unrestricted access to the ICS. Interestingly, even if the attacker had not had access to the firewall configuration files, the device would still have been vulnerable to brute-force attack, permitting "rapid authentication attempts within a short timeframe" that could reveal the administrative password, giving the attacker the ability to use a web interface to manipulate the firewall configuration and gain access to the trusted network (ICS-CERT 2012).

Firewall configuration changes were possible but unnecessary due to the 'dual-homed' nature of the victim. Dual-homed systems involve a single system with multiple network interfaces, each of which is connected to a different network. Dual-homed systems may be the result of unauthorized network connections, but they are also often implemented as part of a system's overall design. In fact, every server and workstation in the test environment was delivered in a dual-homed configuration; each system was simultaneously connected to two different Ethernet networks. Although a dual-homed system does not necessarily constitute the presence of security vulnerabilities, it may provide attackers with the means to gain access to otherwise inaccessible networks. In the test environment, the presence of a single dual-homed system connected to both the internal and external networks provided a foothold that the attacker was able to use to 'pivot' to the internal ICS network.

The cyber-mugging scenario, while unsophisticated, demonstrates the ease with which security perimeters may be breached. Firewalls are essential security devices. However, there are multiple avenues that may take an attacker through or around a firewall. In this instance, the attacker quickly discovered three different methods of fully negating the firewall's security capabilities: by attacking it directly (brute-force), indirectly (configuration utility on workstation), or completely bypassing it (pivot).

Taking control of the internal network

Using Metasploit's built-in routing functionality, the attacker configured Metasploit to take all traffic for the internal subnet and pass it through the Meterpreter session already running on the exploited workstation. Routing traffic to the internal subnet through the compromised host established a pivot point, bypassed the firewall, and facilitated access to the internal network for further attacks.

The attacker initiated reconnaissance of the internal network by using a Metasploit auxiliary module to conduct an ARP scan, revealing all systems and devices on the internal network that were communicating with an IP address. With this new information, a near-complete diagram of the internal ICS network was created and the attacker developed a list of new targets to exploit.

As part of the initial attack, Windows' password hashes were extracted from the external workstation. These hashes were utilized as part of a 'pass-the-hash' attack against the remaining Windows systems. This technique takes advantage of Microsoft's implementation of the NTLM hashing mechanism, in which passwords are not salted when the hash is generated, resulting in a password hash that never changes. The password hash for 'Administrator' on one system will be the same on every system and can be used to authenticate a user throughout the Windows environment. The plaintext password itself is unnecessary, and the complexity of a password is irrelevant (Kennedy *et al*. 2011, pp. 84-85).

To execute the pass-the-hash attack, each Windows host on the internal network was targeted with Metasploit's psexec module, which was configured with the credentials obtained from the compromised workstation. One-by-one, the attacker 'passed-the-hash' and opened Meterpreter sessions with full SYSTEM-level control on all internal Windows systems, giving the attacker full control of the servers and workstations responsible for managing the ICS process and devices.

Range of potential malicious activities

Effectively, any action available to a system administrator was also available to the attacker. Potential malicious actions include shutting down systems; deleting important files; logging operator keystrokes; and sniffing network traffic to reveal plaintext communications, service details, and challenge-response traffic that may be used in offline password attacks (Higgins 2013). The MSSQL database contained user-account details for the ICS, which the attacker modified by adding a new privileged user to the system. The database was also the unexpected location of a table that listed processes to execute when the system starts, an ideal place to hide 'autorun' malware used for persistence.

In addition, the presence of the Siemens engineering/programming tools on one of the servers made it easy for the attacker to conduct a wide-range of ICS-related actions, even without any ICS knowledge. For example, the attacker could use a GUI application to modify PLC control logic to allow performance of unsafe actions or modify functions of existing controls to alter performance in a manner not originally intended or expected. 'Reprogramming' a PLC with no prior experience may seem complicated, but the GUI tools made it easy to point, click, modify values, save the project, and download it to the PLC. The attacker may also inadvertently modify the PLC programming in a way that renders the system inoperable or temporarily 'bricks' the device.

Using functionality provided by the ICS software WinCC Explorer, the attacker had full control over management of the ICS users. The application provided a GUI interface that allows the attacker to easily add unauthorized users, remove or modify the permissions of existing users, modify or reset credentials, lock out users, or perform malicious actions using the account of a legitimate user in an attempt at misdirection. It was also trivial for the attacker to use the ICS

software to modify the HMI user interface elements, perhaps hiding malicious activities from legitimate operators.

Using the Internet Explorer web browser on the compromised systems, the attacker was able to access web applications that were available only by systems on the internal network. Web applications often provide a large attack surface and tend to introduce additional vulnerabilities in to an environment (Bird & Marico 2013). During the exercise, the attacker discovered an SQL injection vulnerability in the Siemens WinCC Web Navigator application running on an internal host. This was not a publically known vulnerability at the time of the exercise, demonstrating that even fully-patched systems have unknown vulnerabilities. This specific vulnerability allowed unauthorized disclosure and modification of data and was publically disclosed in June 2013 as CVE-2013-3957 (US-CERT/NIST 2013).

In addition to web applications, many other network services may be found running on PLCs and network devices in ICS networks. These devices frequently implement TELNET, FTP, SNMP, and other services (Langner 2012) which, in addition to using insecure plaintext protocols for communication, may be enabled by default with insecure configurations. In the test environment, the attacker noticed frequent Dynamic Host Control Protocol (DHCP) requests from Siemens devices requesting IP addresses. The attacker enabled a DHCP server which allowed the devices to automatically obtain IP addresses. Establishing a connection to these newly addressed devices, the attacker discovered that he had access to the Siemens SCALANCE X-series network switches that comprised the communications backbone of the ICS network. Making modifications to the ICS network configuration, such as removing logical segmentation enforced via VLANs, would have been as easy as logging in with the default credentials.

‘Mugging’ the Industrial Control System

Using only the graphical HMI, the attacker was able to control the components of the ICS in a way that might result in physical destruction of plant components. The attacker could open and close physical breakers and valves; turn on pumps; change values of safety and interlock parameters so that equipment operated outside of normal ranges; remove limits on min/max values for parameters such as temperature, fluid volume, and speed; reverse direction of fluid pumps; or allow pumps to run with no fluid in the pipes.

The attacker started by using the point-and-click HMI to open a valve that controls the flow of chemicals into a storage tank. To bypass software-coded safety measures, the attacker modified ‘interlock’ restrictions using functionality available in the HMI. Bypassing interlock restrictions permits execution of a command without the system meeting the programmed interlock conditions. This may allow the attacker to issue commands which would otherwise be prevented for safety reasons, such as keeping a valve open even if the destination tank is already full.

With the interlock conditions bypassed, the attacker opened a valve which allowed chemicals to fill the tank. As the contents of the tank reached dangerous levels, software alarms would have notified other system operators. To prevent alarms from triggering, the attacker changed the alarm parameters to values that would not be met during the attack.

Not content with merely filling the tank, the attacker engaged a pump and steam valve to cycle the tank's contents through a heating system before returning to the tank. Parameters for the steam valve were set to 100% open. Chemicals were pumped from the tank, through pipes leading to the steam heating system, and back in to the original tank. Meanwhile, the tank continued to fill with chemicals from the initially opened valve. As time progressed, the tank reached its maximum capacity while the contents continued to be heated.

Hypothetically, these conditions would result in a failure of physical components of the system; a likely scenario is that the increasing pressure in the tank would force a mechanical safety release valve to open, releasing the toxic chemicals into an emergency storage area. The heating of the chemicals might contribute to aerial dispersal in the form of steam. Additionally, the chemicals in the simulated environment were highly volatile, and explosion would have been feasible. In the event that chemicals were to combust, the resulting explosion would cause significant destruction, and toxic smoke could force the evacuation of surrounding areas. The worst-case scenario would involve violent explosion of the over-pressurized and overheated ammonium nitrate storage tank, potentially triggering secondary explosions throughout the plant, followed by fire and widespread airborne dispersal of toxic fumes.

Comments on mechanical safety devices

Mechanical safety measures should be in place to limit the physical impact of an ICS-related failure. However, these safeties should not be relied upon to completely eliminate the physical risks of a cyber attack. Mechanical safety devices may be poorly maintained, defective, improperly installed, and intentionally bypassed or circumvented for operational purposes. Also, other unforeseen conditions, such as the unpredictable 'human element', may impair the effective operation of a physical safety mechanism. It is not possible to account for all of the variables that factor in to determining whether mechanical devices will function properly. Stated bluntly, things break; if things did not break, accidents would not happen.

In addition, even if machinery fails 'safe' and a life-endangering disaster is averted, the process controlled by the ICS will be negatively affected. In the simulated research environment, a pressure release valve might have prevented over-pressurization of the storage tank, or a mechanical interlock might have closed the steam valve before temperatures reached unsafe levels. While catastrophic failure might have been avoided, the process would have been disrupted or halted. In systems designed for high-availability or with no tolerance for latency, an attack that results in a simple denial-of-service condition might be unacceptable and costly.

Notes on attack-mitigation strategies

The simulated attack documented in this report is one very basic example of many possible scenarios which might ultimately result in the compromise of an ICS environment. It would be counterproductive to pick and choose specific aspects of the attack and assume that, had conditions been slightly different, the attack would have failed. Certain mitigations will absolutely be effective in countering individual techniques illustrated in this exercise; however, there are a multitude of methods for gaining access to internal ICS systems, and there is no panacea or remediation strategy that fully eliminates the risk of an opportunist attack, much less a targeted assault by a sophisticated adversary.

From client-side attacks such as spear phishing, to local/physical access attacks, to thumb drive drops, to the unavoidable presence of 0-day vulnerabilities, the possibility of cracking the security perimeter and obtaining access to ICS networks is always present and, in many instances, trivial. Given sufficient resources, even fully isolated 'air-gapped' systems may be compromised, as demonstrated by the Stuxnet incident (Langner 2012, p. 128). To draw on the world of competitive fighting for an analogy, the defender may attempt to dodge or block every attack, but he or she must also train and fight with the understanding that not all attacks can be evaded. Eventually, the attacker will breach defenses, and the defender must be prepared to take a punch and continue the fight.

Evaluating the effectiveness of mitigation strategies intended to reduce the likelihood of an ICS-related cyber incident was outside the scope of this exercise. However, acknowledging that significant risk is associated with ICS-directed attacks from IT-based attackers, system owners should seek to implement layered defense-in-depth strategies, such as those outlined in *NIST SP 800-82* (Stouffer, Falco & Scarfone 2011) and *ISA/IEC-62443*, formerly *ISA-99*. These types of defensive strategies will raise the level of sophistication required for an attacker to gain unauthorized access to systems and provide defenders with the means to detect the occurrence, mitigate the effects, and recover from a cyber-mugging.

Conclusion

In the event that a malicious actor with limited ICS knowledge gained access to an ICS/SCADA environment, what types of actions are possible, and how disruptive or destructive are the potential consequences? The simulated attack described in this document illustrates that an attacker with a low to moderate level of technical proficiency might obtain full control of the processes and devices in an ICS by utilizing the same publically available tools and basic skills used to compromise a traditional IT domain.

As technological capabilities continue to advance while simultaneously becoming less costly (Weiss 2010, pp. 25-27), the boundary that might once have clearly separated 'control' systems from 'information' systems has eroded. ICS rely heavily on information systems and are, therefore, vulnerable to the same exploitation tools and techniques that are used to target traditional information systems. The simulated process in the test environment was controlled by PLCs, which were managed by software applications running on information systems. By compromising these information systems, an attacker was able to exploit existing trust relationships between systems and devices to manipulate the ICS without any specialized expertise, techniques, or tools. Access to an HMI or the ICS engineering software may further reduce the level of skill required to execute a destructive attack to that of 'point-and-click'.

ICS often rely on IT yet rarely implement IT-security best practices. There are billions of networked control devices; and if only a tiny percentage of these devices are vulnerable to a physically-destructive attack, there are still hundreds of thousands of potential targets. The skills needed to exploit these systems are commonplace within the information-security industry, and the knowledge and tools needed to cultivate these skills are freely available on the Internet. New vulnerabilities are discovered daily, yet ICS rarely receive patches or security updates (estimates indicate that only 10-20% of system owners apply vendor-released patches (Higgins 2013)),

leaving them vulnerable to exploitation. These conditions have produced a target-rich environment for aspiring attackers seeking to create physical destruction through cyber-means.

The potential consequences of destructive ICS incidents may range from administratively annoying to fatally catastrophic. The cyber-physical relationship between information systems and industrial machinery in ICS has exposed environments where even limited resources may be leveraged to trigger significant physical effects. The feasibility that such an incident could cause significant disruptive physical effects directly challenges the current paradigm that state-level resources are required to inflict catastrophic results.

Final thoughts on opportunistic ICS attacks

The attack summarized in this document assumes no ICS/SCADA knowledge on the part of the attacker. While there are significant differences between ICS and traditional IT systems (Weiss 2010, pp. 29-41; Zhu, Joseph, & Sastry 2011), there is also a large degree of overlap in the technologies employed by both types of systems. Furthermore, security considerations and strategies that focus solely on differences between ICS and IT may be irrelevant in situations where ICS are fully-managed by IT systems. An attacker does not need ICS expertise to manipulate ICS/SCADA applications running on a compromised IT system; the IT-based attacker is in control of the system(s) controlling the ICS. As demonstrated by the recent surge in ICS-related security research being presented at conferences, a competent and dedicated attacker with an information-security background could reference open-source information to independently develop *advanced* ICS/SCADA exploitation abilities with only a few months of effort. ICS-tailored security requirements must be implemented in addition to standard IT best practices, not in place of them.

Finally, the attacker's target selection was limited to a single test environment. In a 'real-world' scenario, an opportunistic attacker would begin by seeking a vulnerable target, rather than trying to identify vulnerabilities in a specific target. Services and tools such as Shodan (www.shodanhq.org) and masscan (github.com/robertdavidgraham/masscan) provide the means to quickly discover known vulnerable devices, easily identifying systems as possible 'low-hanging fruit', ripe for attack. Although seemingly countless ICS devices may be discovered with these tools, their effectiveness is limited to exposing systems that are directly accessible via the Internet. It is likely that an exponentially greater quantity of systems and devices remain hidden behind poorly configured firewalls or half buried within global corporate networks.

References

Bird, J & Marico, J 2013, *Attack surface analysis cheat sheet OWASP*, viewed 1 February 2014, < https://www.owasp.org/index.php/Attack_Surface_Analysis_Cheat_Sheet>.

Duggan, DP, Berg, M, Dillinger, J & Stamp, J 2005, *Penetration testing of Industrial Control Systems*, Sandia National Laboratories, Albuquerque, New Mexico, United States.

Higgins, KJ 2013, *SCADA password cracking tool for Siemens S7 PLCs released*, viewed 1 February 2014, < http://www.darkreading.com/vulnerability/scada-password-cracking-tool-for-siemens/240146748>.

ICS-CERT 2012, *Siemens SCALANCE S Multiple Security Vulnerabilities|ICS-CERT*, viewed 1 February 2014, < http://ics-cert.us-cert.gov/advisories/ICSA-12-102-05>.

—2013. *Siemens SCALANCE Privilege Escalation Vulnerabilities|ICS-CERT*, viewed 1 February 2014, <https://ics-cert.us-cert-gov/advisories/ICSA-13-149-01>.

Kennedy, D, O'Gorman, J, Kearns, D & Aharoni, M 2011, *Metasploit: the penetration tester's guide*, No Starch Press, Inc., San Francisco, California, United States.

Langner, R 2012, *Robust control system networks: how to achieve reliable control after Stuxnet*, Momentum Press, LLC, New York, New York, United States.

Microsoft Corporation 2008, *MS08-067: Vulnerability in server service could allow remote code execution*, viewed 1 February 2014, <http://support.microsoft.com/kb/958644>.

National Institute of Standards and Technology 2013, *Security and Privacy Controls for Federal Information Systems, NIST Special Publication 800-53 revision 4*, National Institute of Standards and Technology, Gaithersburg, Maryland, United States.

Siemens AG 2005, *SIMATIC NET SCALANCE S and SOFTNET security client operating instructions.*

Stouffer, K, Falco, J & Scarfone, K 2011, *Guide to Industrial Control Systems (ICS) Security, Special Publication 800-82*, National Institute of Standards and Technology, Gaithersburg, Maryland, United States.

US-CERT/NIST 2013, *National Vulnerability Database (NVD) vulnerability summary for CVE-2013-3957*, viewed 1 February 2014, <http://web.nvd.nist.gov/view/vuln/detail?vulnID=CVE-2013-3957>.

Weiss, J 2010, *Protecting Industrial Control Systems from electronic threats*, Momentum Press, LLC, New York, New York, United States.

Zhu, B, Joseph, A & Sastry, S 2011. *A taxonomy of cyber attacks on SCADA systems, 31 March 2014*, <http://bnrg.cs.berkeley.edu/~adj/publications/paper-files/ZhuJosephSastry_SCADA_Attack_Taxonomy_FinalV.pdf>.

Building Future Generations of Elite Cyber Professionals (CNODP)

NSA's Computer Network Operations Development Program Staff

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: cnodp_all@nsa.gov

Abstract: *With the increase in cyber attacks, defending America's networks is one of the primary Department of Defense challenges in the 21st century. It is a national imperative to have elite cyber-warfare forces trained and ready to protect the country's National Security Systems and critical infrastructure against attacks in cyberspace. To that end, the National Security Agency has created the Computer Network Operations Development Program (CNODP), a highly effective cyber-defense workforce-training program. The CNODP is NSA's premier vehicle for developing skilled civilian and military personnel into highly effective cyber warriors and capability creators who build on their degrees in computer science, electrical and computer engineering, mathematics, and information assurance. Rotational work assignments provide program participants with challenging technical experiences in multiple locations, missions, and disciplines, as well as continual and enduring networking and mentorship within the broader Computer Network Operations community.*

Keywords*: NSA; DoD; Cyber; CDX; Exercise; Red Team; Defense; Attack; Exploitation; Protect; Training; Computer Network Operations Development Program; Software, Hardware; Knowledge, Skills, Attitudes (KSA); Programming; Reverse Engineer; Rotation; Rotational Assignments; SELinux*

Introduction

As the challenges of cyber security increases so does the demand for subject matter experts in Computer Network Operations (CNO). The Department of Defense needs trained cyber warriors who can thwart cyber attacks and avoid what former Defense Secretary Leon E. Panetta calls a '"cyber-Pearl Harbor that would cause physical destruction and the loss of life, an attack that would paralyze and shock the nation and create a profound new sense of vulnerability'" (Bumiller & Shanker 2012). The NSA welcomes President Obama's government-wide review of federal job training programs and is confident the Computer Network Operations Development Program (CNODP) will prove to be not only one of the most successful ones, but also one that warrants emulation. Graduates of the CNODP are highly sought after; organizations within and outside of NSA compete to persuade these graduates to become members of their teams.

What Is Computer Network Operations (CNO)?

Broadly speaking, Computer Network Operations consists of Computer Network Attack (CNA), Computer Network Exploitation (CNE), and Computer Network Defense (CND). CNA includes actions taken via computer networks to disrupt, deny, degrade, or destroy the information within

computers and computer networks and/or the computers/networks themselves. CNE involves enabling operations to gather data from target or adversary Automated Information Systems (AIS) or networks. CND includes using computer networks to analyze, detect, monitor, and protect against attacks, intrusions, disruptions, or unauthorized access to the network. Defending the United States' networks against attacks is a major concern. Indeed, former NSA Director Vice Admiral John McConnell has said the U.S. is ""more vulnerable than any nation on earth"" (Washington 1995).

NSA's Computer Network Operations Development Program (CNODP)

NSA created CNODP in 1995 in response to a growing need to develop the existing technical expertise at NSA in the areas of Computer Network Operations to meet future requirements. The demand for cyber experts who can stay ahead of cyber threats is increasing and will continue to grow. As Sonja Treven, Associate Professor at the School of Business and Economics at the University of Maribor in Maribor, Slovenia, explains:

> The security of computer networks will continue to increase in importance as more business is conducted over the Internet. Organizations need to understand how their systems are vulnerable and how to protect their infrastructure and Internet sites from hackers, viruses, and other acts of cyber-terrorism. (2004)

The program was initially envisioned to be a new-hire program, but it became a retraining program for highly qualified military and civilian employees working at NSA. Today, the program consists of several dozen new hires, as well as military and NSA civilian employees. The duration is three years, so there are over 100 participants in the program at any given time. Participants in the program are called interns because they are training and learning on the job how to apply their technical skills and education in the Computer Network Operations field. They learn from the best by working beside elite cyber experts who are knowledgeable about the latest cutting-edge technology.

As Steve LaFountain, one of NSA's technical leaders, has stated, ""The nation increasingly needs professionals with highly technical cyber skills to help keep America safe today and to help the country meet future challenges and adapt with greater agility"" (NSA 2012). The CNODP seeks to meet this need by attracting employees who are curious about how things work and love to solve problems. Obtaining a position in the program is competitive; only the best and the brightest are selected. The positions advertised are for entry-level employees or those who have up to five years of experience.

The goal of CNODP is to educate and train cyber warriors in the various technical areas of Computer Network Operations. The three-year, individual-development program begins with three months of challenging training in the classroom with hands-on laboratory experience. Immediately following the training, employees spend the next thirty months completing rotational-tour assignments during which they have the opportunity to apply and grow their technical skills in software development, networking, or hardware. In the last three months in the program, employees work together on a capstone project that solves one of NSA's hard problems.

The Benefits of CNODP

Guidance from expert mentors and advisors

Mentors are essential and play a key role in providing guidance and support. The interns must choose a mentor for the duration of their three years in the program. All of the mentors available to the interns are members of the CNO community, and many of them are CNODP graduates. In collaboration with his or her mentor, each intern must work to complete an Individual Development Plan (IDP), which maps out the person's goals while in the program, current strengths, focus areas for improvement, and rotational assignments. The IDP can be changed as often as needed, but major revisions must be approved by the mentor. The plan may also change if an office realigns or has limits on how many interns it can support at one time. In addition, interns are supported through the program's curriculum by expert advisors. One class of interns consists of 40 individuals; and there are four advisors who follow each class of interns through its entire three-year program. The advisors review all of the IDP's for the entire class to ensure they follow the CNODP program philosophy. The advisors also meet with members of the class annually for general feedback and guidance. In addition, advisors are heavily involved with the class project during the last three months of the three-year program.

Core training

Interns in CNODP start the three-year program by completing three months of in-house training at NSA that is taught at the master's level. The training is labeled 'core' because it is the central (main or most important) component that provides the basic knowledge the interns need to be successful in the program. Since the interns have a wide range of backgrounds, the training is a workforce-development plan that addresses the basic Knowledge, Skills, and Abilities (KSA) that are applicable to the CNO tasks performed in NSA organizations that host interns in rotational assignments. Because there are a growing number of colleges and universities that offer cyber-security courses, some interns join CNODP with many of the required KSAs for the program. NSA has designated eight schools as national Centers of Academic Excellence (CAE) in Cyber Operations in an effort to broaden the pipeline of cyber warriors.

The CAE-Cyber Operations Program is intended to identify institutions offering a curriculum that is deeply technical, interdisciplinary, and firmly grounded in computer science, computer engineering and/or electrical engineering, with extensive opportunities for hands-on applications via labs/exercises (Dodge 2010). This program is in support of the president's National Initiative for Cybersecurity Education (NICE): Building a Digital Nation. The two primary objectives of this initiative are to provide specialized training to the federal cyber workforce and to ensure a federal cyber-security workforce pipeline for the future. As an additional objective, NICE also focuses on protecting cyberspace. The CNODP has benefitted from NICE and hired numerous interns from the CAE Cyber Operations schools. The interns hired from CAE Cyber Operations schools do extremely well in the three months of core training that are held at NSA.

CNODP interns are trained by instructors who are subject matter experts, members of the CNO community, and masters of the technical skills in the area they are teaching. The curriculum presupposes prior experience and is both demanding and comprehensive. For example, interns must be proficient in the programming language C since the first course in core training is essential C programming, and additional core training courses utilize C programming. Operating systems (OS), microcontrollers, and many systems are written in C, and students can develop programs that can be embedded into an OS kernel, such as a device driver, with the required complexity and sophistication to implement exploits for discovered vulnerabilities. A Python programming course is a fairly new addition to core training, and it is used as a scripting

language that allows versatility such as fewer lines of code than C. Windows and Linux OS courses apply the student's knowledge of C programming language, and interns are also taught how to use CNO reverse engineering tools.

Core training is evaluated annually, and the appropriate changes are made in accordance with advances in tools and technology. The CNODP Implementation Board directs the generation of processes and procedures for recommendation to the program offices; so the board has approval authority over changes to core training.

Educational and professional enrichment opportunities during core training

Members of the CNO community know all too well how important it is to engage in intensive continuing education:

> Employers and employees have struggled to keep pace with change, making the development of a formalized career model very challenging. Because specific job roles will shift with the advent of new threats and new technologies…competency in core skills is essential. These capabilities include both quantitative skills such as engineering, mathematics and computer science, as well as behavioral skills such as management, communication and the ability to think creatively. Thus, the demand for cybersecurity expertise cannot easily be described with a uniform skill profile. Rather, needed expertise encompasses an ecosystem of complementary knowledge, skills and abilities. (Dodge 2010)

Because the NSA understands this imperative, the CNODP provides a variety of opportunities for the interns to engage in additional educational opportunities during their core training. Of course, the interns have occasions to network and help each other; this fosters teaming which is a strong development goal of the program. In addition, however, interns can participate in NSA's After-Hours College Program, which allows employees to take undergraduate or graduate courses that relate to NSA's mission at accredited colleges or universities, with NSA paying for the tuition and lab fees. Interns may also work towards a Master of Science degree by taking courses during their work day, with supervisor approval, at the Naval Postgraduate School (NPS), which is one of the eight CAE Cyber Operations Schools. Moreover, because each organization has its own training requirements that necessitate additional coursework or self-study, interns also have the opportunity to grow by attending additional training within or outside NSA, such as conferences and classes, in order to meet their specific goals.

Rotational assignments or touring

Of course, after completing core training, interns are anxious to apply their skills to real-world problems; and because of their extensive training, they are ready to perform 30 months of rotational assignments. Touring in offices broadens their knowledge of CNO. It is a great way to motivate employees, and it pushes them to do their best. Rotation allows interns to work in different positions at various levels. They can determine what they are good at as well as what they enjoy doing. A stretch tour is a tour in an office whose primary functions are ones that are not immediately familiar to the intern or part of his or her technical skills and background. For instance, an intern with software development skills may tour in an office that performs hardware functions. The tour allows the intern to perform evaluations of microelectronic devices in order to understand if vulnerabilities exist and what impact they may have on the functionality of the

device. The intern can also work in a laboratory to perform physical processing of integrated circuit to understand architecture, electronic circuit routing, and physical features of integrated circuit. Because the NSA values this type of experience, stretch tours are highly encouraged in the program.

Summer Conference on Applied Mathematical Problems (SCAMP) option

For ten weeks each year, interns have the option to diversify from rotational assignments and participate in a Summer Conference on Applied Mathematical Problems (SCAMP). They work together with some of the best computer scientists, mathematicians, and statisticians on challenging problems relevant to national security. SCAMP is held at The Institute for Defense Analyses Center for Computing Sciences (IDA/CCS) in Bowie, Maryland, which is an independent, applied research center. Most of the faculty members have doctorate degrees in mathematics, computer science, and computer and electrical engineering. IDA/CCS performs research in network security and methodology for mining large data sets. At the end of the ten weeks, the interns return to NSA with increased technical knowledge that will be applied to the Agency's CNO mission.

Cyber Defense Exercise (CDX)

Interns also have the opportunity to participate in the Cyber Defense Exercise (CDX), which is a four-day computer-security competition. This Information-Assurance annual event allows participants to tackle cyber threats in a simulated real-world exercise. They perform CND to detect attacks and protect a network of computers, as well as try to keep their network secure and operational:

> "CDX offers an unparalleled opportunity for some of the nation's top students to showcase their cyber skills to NSA's leading practitioners," said Neal Ziring, IAD's Technical Director. "America increasingly needs professionals with highly technical cyber skills to help the country remain safe and adapt with greater agility. We need the best and brightest to help us defeat our adversaries' new ideas". (NSA 2013)

The benefits of this annual competition extend beyond the hands-on educational enrichment of the interns. The lessons the interns and other competitors learn during the Cyber Defense Exercise are shared with government and academia.

CND tours

Certainly some of the most important elements of the program are the required CND tours. Each intern is required to work in CND missions while in the program. NSA's Red Team is one of the most popular CND mission elements interns select for their defensive tour. The Red Team's mission is to identify vulnerabilities and weaknesses in United States' cyber-information systems. The Red Team simulates real-world CNO adversary or opposition forces during DoD and government assessments, exercises, and Information-Operations activities; demonstrate**s** the impact of identified vulnerabilities and weaknesses in a near-real-world environment; and recommends ways to mitigate identified vulnerabilities and weaknesses.

The Red Team is just one option for CND touring. In fact, the interns have the opportunity to apply and grow their cyber skills in various offices at NSA. This gives interns the chance to gain invaluable experience and work with some of the country's best analysts, operators, and

developers. As General Keith B. Alexander, former Director of NSA and U.S. Cyber Command, said, the United States maintains a

> "deep, persistent and pervasive presence on adversary [computer] networks. We maintain that access, gain deep understanding of the adversary, and develop offensive capabilities through the advanced skills and tradecraft of our analysts, operators and developers. When authorized to deliver offensive cyber effects, our technological and operational superiority delivers unparalleled effects against our adversaries' systems". (Waterman 2013)

CNODP interns not only observe, but also participate in the development of these capabilities. For example, one of our CNODP interns joined the SE Android team in the Research Directorate during one of his tours. This team's mandate is to enhance Android security by adding Mandatory Access Controls (MAC), similar to the work done a decade ago to create SELinux. After researching the weaknesses of the Android permission model, he found that programmers, even engineers from handset manufacturers, barely understood Android's inter-app communication model. After learning about the threats faced by Android middleware, the intern began to devise a solution to prevent unauthorized inter-app communications by requiring explicit authorization for communications. As a result, a prototype was released on the public SE for the Android Git repository, which sparked some interest from external researchers and handset manufacturers. Clearly, the benefits of these CND tours extend well beyond the experience of the interns who take part in them.

Capstone project and placement

The three-year program ends with a three-month capstone project. The project focuses on a CNO-related problem that the Agency considers one of its 'hard problems.' NSA organizations actually compete to have interns select their projects. The CNODP Board of Governors has the final decision authority for all aspects of the program, which includes approving the class project. The board consists of high-level decision makers who have the authority and/or access to those who have control over agency resources that affect the program. Once the capstone project is selected, the group of highly skilled interns works together for three months to solve the problem. The project allows the interns to exercise and demonstrate the teamwork, adaptability, flexibility, and drive that are expected of program graduates. After the final project has been completed, interns leave the program and are placed in an organization with a CNO-related mission that has been approved by the Board of Governors. The board's decision factors in the desires of the intern and the needs of the agency. Graduates from the program have become technical leaders, subject matter experts in their technical fields, and driven members of the CNO community.

Conclusion

NSA continues to lead the way in recruiting and training highly skilled cyber warriors. CNODP fulfills President Obama's NICE by providing specialized training to the workforce and utilizing the pipeline of graduating students from CAEs and schools that are producing talented and promising cyber warriors. While in CNODP, interns receive a diverse range of exposure in the core areas of software, networking, and hardware. They graduate from the program with the ability to solve an important problem, independent of the technology at the core of the problem.

Technology is constantly changing, and NSA is adapting to the changes as well as offering

training in the newest technologies. There is a growing network of CNO professionals at NSA who support its mission by teaching CNO-related courses, providing mentorship, and applying their technical expertise. This professional support network is essential to maintaining a thriving innovative culture in our CNO workforce and a critical element in improving national security systems.

The adversary is using advanced technology and may be capable of launching multiple attacks on critical U.S. infrastructure. The threats from cyber attacks are real and escalating. Because CNODP is a program that is successfully creating and training cyber warriors to meet that challenge, it is an initiative that warrants emulation throughout the DoD.

References

Bumiller, E & Shanker, T 2012, 'Panetta warns of dire threat of cyberattack on US', *New York Times* 11 October, <http://www.nytimes.com/2012/10/12/world/panetta-warns-of-dire-threat-of-cyberattack.html ?pagewanted=all&_r=0>.

Dodge, RC, Toregas, C, & Hoffman, L 2012, 'Cybersecurity workforce development directions', *Proceedings of the sixth international symposium on human aspects of information security and assurance*, pdf, Springer-Verlag, Heidelberg, <http://www.cspri.seas.gwu.edu/uploads/2/1/3/2/21324690/costis_-_cybersecurity_workforce_development_directions.pdf>.

NSA Public and Media Affairs 2012, press release, NSA, 21 May, <http://www.nsa.gov/public_info/press_room/2012/new_college_cyber_ops_program.shtml>.

—2013, press release, NSA, 19 April, <http://www.nsa.gov/public_info/press_room/2013/cdx_2013.shtml>.

Treven, S 2004, *The need for training of computer and information systems managers*, pdf, viewed 15 May 2005, p. 91, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.182.2911&rep=rep1 &type=pdf .

Washington, DW 1995,'Onward cyber soldiers', *Time Magazine* 21 August, vol. 14, no. 8, 20140317, <http://www.csm.ornl.gov /~dunigan/timemag.html>.

Waterman, S 2013. 'U.S. cyberwar offense "best in the world": NSA's Gen. Keith Alexander', *The Washington Times*, 26 August, <http://www.washingtontimes.com/news/20 *13/* aug/26/us-cyberwar-offense- best - world-nsas-gen-keith-alex/#ixzz2reYoHSEC>.

Introducing the National Security Cyber Assistance Program (NSCAP)

G Hale, R Lenzner

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

***Abstract**: The demand to improve the robustness and survivability of National Security Systems (NSS) continues unabated, as each year the number and sophistication of cyber attacks from a variety of sources—nations, criminal and hacking groups, and individuals with malicious intent—increase. As a measure to address this problem, the National Security Agency/Information Assurance Directorate (NSA/IAD) has launched the NSCAP accreditation initiative, designed to make available a pool of qualified cyber-security service providers for supporting NSS owners and operators when similar services are not readily available from within their organizations. The NSCAP accredits service providers who meet NSA-developed criteria in the field of Cyber Incident Response Assistance (CIRA). In the future,NSCAP may be expanded to include the accreditation of Vulnerability Assessment and Penetration Testing (VAPT) service providers.*

Keywords: *Incident Response Assistance, Vulnerability Assessment and Penetration Testing, NIST 800-53 IR-7, ISO 27001, CNSSP 22, CNSSI No. 1253*

Introduction

Government and industry experts recognize that the need to manage and protect Information Technology (IT) resources and the information that they contain is becoming more critical (Microsoft 2013, pp. 5-8). The growing sophistication and number of cyber attacks necessitate the effective implementation of secure architectures and designs, verification through real-world testing, and an equally vigorous and effective rapid-response capability from system owners and operators when an incident occurs. A primary concern to these owners and operators is that malicious cyber actors are becoming more persistent, capable, and efficient in exploiting a system's vulnerabilities, both those known and recognized, as well as those newly discovered (APT1 2013, p. 2).

As established by federal policies (The White House 1990), which outline the roles and responsibilities for securing National Security Systems (NSS), and applicable sections of the Federal Information Security Management Act of 2002 (Title III of the E-Government Act 2002), a primary function of the NSA/IAD is to provide guidance and support to owners and operators of IT systems within the Department of Defense (DoD) and Intelligence Community (IC), with a focus on systems designated as NSS. NSA/IAD's guidance and support to NSS owners and operators fall into two broad categories:

- Proactively defending systems by performing security testing and providing instructions on the implementation of security measures, delivered through Vulnerability Assessments and Penetration Tests (VAPT) (NIST 2008), and
- Providing reactive Cyber Incident Response Assistance (CIRA) in identifying, containing, and remediating the most advanced and persistent threats and compromises (NIST 2013). CIRA includes both advanced Intrusion Detection and Incident Response services provided as assistance to an NSS owner's internal incident response capabilities.

What Is a National Security System?

- A National Security System is *any information system (including any telecommunications system) used or operated by an agency, a contractor of an agency, or other organization on behalf of an agency where the use or operation involves intelligence activities, cryptologic activities related to national security, command and control of military forces, or is an integral part of a weapon or weapons system*. This definition also encompasses *systems used to support military or intelligence missions*.
- *Systems that are specifically identified within approved legislation or executive order and must be kept classified in the interest of national defense or foreign policy* are also considered to be NSS.
- *All U.S. government classified networks* have been designated as NSS.
- The DoD has included *the Non-Secure Internet Protocol Router Network (NIPRNet)* as an NSS. (NIST 2003; CNSSI 4009 2010)

To improve effectiveness and expand capacity to meet NSS owners' requests for CIRA and VAPT services, NSA/IAD established the NSCAP. The program intends to 1) develop a list of exceptionally capable and highly qualified cyber-service providers that meet the minimum accreditation requirements that are available to the NSS community, 2) promote public-private collaboration, 3) leverage industry expertise to protect national interests, and 4) improve the U.S. Government's ability to address and manage current and emerging cyber threats.

Although CIRA and VAPT accreditation frameworks are grouped under NSCAP, processes followed to perform these services differ in one key aspect. Time scales weigh heavily into the level of engagement and planning, with those for CIRA being vastly shorter than those for VAPT services (Lewis 2010). Since VAPT services are pre-planned and scheduled in advance, the service provider can perform detailed planning prior to delivery to the client. By contrast, the need for CIRA services is usually urgent; therefore, technical details of the engagement scope are usually less than comprehensive.

As outlined in the following sections, NSA/IAD will follow a phased approach to administer CIRA and VAPT accreditation frameworks under this structure, first addressing the requirements for CIRA accreditation. NSA/IAD recently completed the pilot phase of the CIRA accreditation program and is transitioning this program to a full production state. After a successful CIRA roll-out, NSA/IAD will complete the development of the VAPT accreditation and phase it into the program. This approach will allow for the development process to draw from positive aspects of CIRA accreditation through the utilization of lessons learned.

CIRA Accreditation

The core objective of CIRA accreditation is to identify companies qualified to provide rapid, on-site support to NSS owners and operators in incident response and intrusion detection. Broadly speaking, assessment of capabilities is based on the provider's ability to

- Consistently deliver services using repeatable processes and procedures;
- Assign highly skilled and qualified staff, who are eligible to hold U.S. Government security clearances, to follow outlined processes and procedures to deliver services; and
- Maintain and improve the quality of delivered services through training initiatives, improvement of analytical capabilities, and use of lessons learned from previous deployments or engagements to refine processes.

NSA/IAD used survey data, interviews, and research to deconstruct the CIRA framework into separate phases, assess the feasibility of incorporating features from existing accreditation programs into the CIRA model, and identify published standards and guidelines relevant to CIRA certification. The methodology consisted of the following actions:

1. Identify CIRA services offered by DoD and federal civil sector Cyber Incident Response (CIR) Teams/Cyber Emergency Response Teams,
2. Identify and review CIRA services offered by industry,
3. Examine CIR commercial certifications available to individuals,
4. Examine CIR certification programs offered by other U.S. Government organizations and by the British Government Headquarters Communications (GCHQ), and
5. Examine U.S. Department of Commerce, National Institute of Standards and Technology (NIST), and International Standards Organization (ISO) CIR Requirements.

Drawing from the results of steps 1 and 2, NSA/IAD subdivided CIRA services into four phases.

1. Engagement Agreement and Planning: The NSS owner/operator and service provider agree on the scope of the work and then plan how to implement the engagement agreement.
2. Detection and Analysis: The service provider, using both client and the company's own resources, identifies the source of the incident, analyzes the threat, performs an initial damage assessment, and provides containment recommendations. Damage assessments and containment recommendations are derived from the company's performance of complex log analysis, host integrity checking, and network traffic flow analysis.
3. Containment and Remediation: The service provider assists the client in implementing the agreed-upon containment plan and provides support for system remediation.
4. Post-Incident Reporting and Lessons Learned: The service provider presents a recap and out-brief to the client. The provider must also conduct an internal lessons learned of what worked within the engagement, identify issues, and develop an internal corrective action plan to ensure that those problems do not occur in the future.

After examining CIR-certification programs offered by other USG organizations, as well as those offered by GCHQ (CESG | Cyber Incident Response 2013), NSA/IAD concluded that developing a program would not sufficiently address operational needs of CIRA. To establish the boundaries of this accreditation, a comprehensive survey of published guidelines and standards applicable to CIRA accreditation was performed. This review included the most recognized cyber-security documents issued by NIST (Special Publications 500 & 800 Series) and the ISO/International

Electrotechnical Commission (IEC) Information Security Management Standards (ISMS), also known as the ISO 27000 series (ISO/IEC JTC 1/SC 27).

After review and discussions, NSA/IAD's incident response team identified twenty-one core areas of expertise that a candidate organization must possess and excel in to earn CIRA accreditation. These are listed in **Table 1**.

CIRA Application Package Content Requirements		
Business Statement of Intent	Service Agreement Criteria	Log Collection and Analysis Criteria
Core Capabilities Overview	Client Engagement Management Process Criteria	Network Traffic Data Collection and Analysis
Process and Procedures	Communication Management Process Criteria	Host Integrity Data Collection and Analysis
Staff Skill Documentation	Preliminary Data Collection Criteria	Incident Analysis
CIRA Education and Training Plan Criteria	Engagement Tools and Resources Identification Criteria	Containment and Remediation Recommendation
Past Performance Criteria	Travel Management Process Criteria	Post-Incident Analysis
Client Furnished Information and Data Criteria	Rules of Engagement Criteria	Lessons Learned

Table 1: CIRA Application Package Content Requirements

These areas of expertise were optimized to reflect the expected needs of the NSS community and presented to industry within the accreditation instruction manual (NSA/IAD 2013). This manual is currently being updated to incorporate lessons learned from the Pilot Program.

VAPT Accreditation

The NSA/IAD VAPT Service Provider accreditation is being developed to meet the growing needs of the U.S. Government by recognizing service providers that are available to supplement, not to replace, internal capabilities of NSS owners and operators in this area. These services, required for the accreditation of NSS under NIST SP 800-53 Rev 4 (2013) and SP 800-115 (2013), include basic vulnerability assessments with limited or no vulnerability confirmation, to full penetration testing with escalation of privileges and services across the target environment.

NSA/IAD deconstructed the provisioning of VAPT services into four phases.

1. Engagement Agreement and Planning: The NSS owner/operator and service provider agree on the scope of the work and then plan how to implement the engagement agreement.
2. Discovery (Vulnerability Assessment): Within the terms of the Engagement Agreement, the service provider confirms and/or identifies the system or target space of interest, then assesses, identifies, and documents likely vulnerabilities.

3. Exploit/Escalation (Attack/Penetration Testing): Within the terms of the Engagement Agreement, the service provider tests the previously identified vulnerabilities to confirm they are exploitable. If the vulnerability is confirmed, it may be exploited with escalation of any system access. On escalation, it may be appropriate for the scope of the agreement to be expanded and Phases 1-3 repeated.
4. Reporting and Recommendations: Once the assessment, testing, and data collection portions of the Engagement Agreement have been satisfied, the service provider prepares a final report, which includes a set of comprehensive remediation recommendations for consideration by the NSS owner/operator.

The following criteria form the basis for evaluating a service provider's ability to consistently deliver VAPT services and achieve accreditation in this area.

- Uses a framework of documented processes and procedures that ensure
  - Thorough planning of engagement-service requirements,
  - Provisioning of assessments to include analysis of deployed security architectures,
  - Common understanding and documented agreement on legal parameters and constraints,
  - Safe and efficient execution of services,
  - Secure handling of collected data,
  - Prompt delivery of final reports,
  - A final report that includes mitigation and remediation recommendations, and
  - Total satisfaction of all engagement objectives as defined by the engagement agreement.
- Assigns highly skilled and qualified staff that
  - Use established processes and procedures as consistent, repeatable starting points for each engagement and
  - Possess the experience yet are flexible and creative enough to develop and provide customized and adaptive solutions to meet objectives identified in the engagement agreement.
- Maintains and improves the quality of VAPT service offerings through
  - Staff training and training plans,
  - Vulnerability identification and exploit research,
  - Development of enhancements to discovery and exploit tool sets, and
  - Use of a formal lessons-learned process to improve the effectiveness of future engagements.
- Assumes responsibility for services provided in accordance with its engagement agreement.
- Employs VAPT staff members eligible to receive a U.S. Government security clearance (applicants are not required to have a security clearance to apply for and receive accreditation).
- Provides evidence of at least three discrete, different, and successful VAPT service engagements.

Way Forward

The NSCAP continues to enhance the CIRA accreditation and anticipates continuing to develop the VAPT service capability requirements to meet the current and evolving needs of NSS owners

and operators. Refinements to the existing CIRA requirements for intrusion detection and incident response are intended to meet the changing NSS operational environment as newer technologies are deployed. In its current framework, CIRA accreditation focuses on performance of intrusion detection and incident response in conventional NSS environments. In general terms, these are environments that contain wired networks, servers, and workstations, both physical and virtual. Future CIRA accreditation requirements may demand that applicants meet established criteria for assessing wireless networks (newer to NSS environments), mobile devices, and cloud computing.

Conclusion

As technologies change, cyber defenders and attackers are continually evolving tactics and methodologies. For computer network defense, the objectives of defenders remain to thwart attacks and defend vital communications and information from malevolent elements seeking to gain profit, deny critical services, steal sensitive information, or commit other actions among a long list of possibilities. The NSA is charged by law and directive to defend and harden U.S. Government NSS.

To strengthen efforts and expand capacity in protecting NSS against a growing threat, the NSA is solidifying its strategic partnership with commercial industry. In 2013, NSA/IAD established the NSCAP as a means to document NSS requirements for cyber incident response assistance, while providing a standardized program for ensuring integrity and accountability of providers responding to requests for assistance. With completion of the initial NSCAP pilot phase, the first CIRA accreditations will be issued to commercial entities that have met NSA/IAD performance criteria. Buoyed by the assurance that identified providers have conformed to established CIRA standards and demonstrated that they have the processes and expertise in place, NSS owners and operators have a choice of accredited partners from whom they can request assistance.

References

APT1 (Advanced Persistence Threat): exposing one of China's cyber espionage units 2013, Mandiant Investigative Report, Mandiant, A FireEye Corporation, Alexandria, Virginia, United States, p. 2.

CESG|Cyber incident response 2013, cyber incident response, viewed 22 January 2014, <www.cesg.gov.uk/service catalogue/cir/pages/Cyber-Incident-Response.aspx>.

CNSS National Information Assurance (IA) glossary, CNSSI 4009 2010, Committee on National Security Systems, CNSS Instruction (CNSSI), Fort Meade, Maryland, United States.

ISO/IEC JTC 1/SC 27 IT Security techniques, International Standards Organization Publication, International Standards Organization (ISO)/International Electrotechnical Commission (IEC), Information Security Management Standards (ISMS), Geneva.

Lewis, E 2010, 'Information security management', in H Bidgoli (ed), *The handbook of technology management*, John Wiley & Sons, Inc., Hoboken, New Jersey, United States, p. 940.

Microsoft 2013, *Software vulnerability exploitation trends*, report, Microsoft Security Engineering Center, Redmond, Washington, United States, pp. 5-8.

NIST Special Publications (500 & 800 Series), Information Technology Laboratory, National Institute of Standards and Technology Computer Security Division, Gaithersburg, Maryland, United States.

NIST 2003, *Guideline for Identifying an Information System as a National Security System. Special Publication 800-59*, August, Special Publications (800 Series), Information Technology Laboratory, National Institute of Standards and Technology Computer Security Division, Gaithersburg, Maryland, United States.

—2013, *Security and Privacy Controls for Federal Information Systems and Organizations. Special Publication 800-53 rev 4 2013, IR family of controls, IR-7 IR Assistance 2013*. Special Publications (800 Series), Information Technology Laboratory, National Institute of Standards and Technology Computer Security Division, Gaithersburg, Maryland, United States, pp. F-108.

—2013, *Security and Privacy Controls for Federal Information Systems and Organizations. Special Publication 800-53 rev 4 2013*, Special Publications (800 Series), Information Technology Laboratory, National Institute of Standards and Technology Computer Security Division, Gaithersburg, Maryland, United States, pp 1-457.

—2008, *Technical Guide to Information Security Testing and Assessment. Special Publication 800-115 2008*, Special Publications (800 Series), Information Technology Laboratory, National Institute of Standards and Technology Computer Security Division, Gaithersburg, Maryland, United States.

NSA/IAD 2013, *The NSA/IAD Cyber Incident Response Assistance (CIRA) Service Accreditation Instruction Manual, Version 2.0*, NSCAP manuals, Remote Deployment Operations, Information Assurance Directorate, Fort Meade, Maryland, United States.

Title III of *the E-Government Act (Public Law 107-347)* 2002, *Federal Information Security Management Act (FISMA)*, 17 December.

The White House 1990, *National Security Directive 42,National Policy for the Security of National Security Telecommunications and Information Systems*, NSD42 Section 5.d, 11.e, 5 July, the White House, Washington, D.C., United States, pp.8-10.

Active Cyber Defense: A Vision for Real-Time Cyber Defense

MJ Herring, KD Willett

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *Cyber operations consist of many functions spanning cyber management, cyber attack, cyber exploitation, and cyber defense, all including activities that are proactive, defensive, and regenerative in nature. A subset of cyber defense, Active Cyber Defense (ACD) focuses on the integration and automation of many services and mechanisms to execute response actions in cyber-relevant time. ACD is comprised of a set of logical functions to capture details from enterprise-level architecture to operational realization with the primary objective to become a living part of DoD cyber operations to help defend the nation from cyber-based adversaries.*

Keywords*:* *Active Cyber Defense (ACD), Cyber Defense (CD), Sensing, Sense-Making, Decision-Making, Message Fabric, Shared Situational Awareness, Automated Response Action, Coordinated Response Action*

Introduction

In their *2013 Data breach investigations report*, Verizon notes that while 24% of the initial compromise stage of data intrusions takes minutes or seconds, the predominant number of initial compromises take hours. These breaches consist of a series of actions performed in real-time that lead to a persistent malicious presence in the targeted network. Per the Verizon report, discovery of malicious activity by network owners is currently on the order of months, meaning that malicious actors have time to exfiltrate terabytes of data and perform other malicious acts that are unlikely to draw attention in a timely manner.

Recognizing the need to accelerate detection and response to malicious network actors, the United States (US) Department of Defense (DoD) has defined a new concept, Active Cyber Defense (ACD) as "DoD's synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities" (DoD 2011). Among the many needs of war-fighter operations, there is the need to be secure, which includes the concepts of hardening, protecting, attacking, and defending among the war-fighter domains of land, sea, air, space, and cyber. Cyber is an integrating capability for the other domains, as well as a standalone domain that has its own unique needs for cyber defense.

Cyber defense includes three complementary categories: 'proactive', 'active', and 'regenerative'. 'Proactive' activities harden the cyber environment and maintain peak efficiency for cyber infrastructure and mission functions. 'Active' activities stop or limit the damage from adversary

cyber activity in cyber-relevant time. 'Regenerative' activities restore mission effectiveness or efficiency after a successful cyber attack. These categories form a continuum of cyber-security activities occurring continuously and simultaneously on networks, integrated by a common framework of automation that includes ACD as a subset of integrated cyber defense. The focus herein is on ACD.

ACD is purposely designed to be applicable across the U.S. Government (USG) as well as Critical Infrastructure and Key Resources (CIKR). This facilitates the reuse of ACD-related solutions across the USG and CIKR. Such leveraging is fiscally responsible to the U.S. taxpayer and ultimately minimizes the total cost of ownership for ACD across the USG. The Information Assurance Directorate's (IAD) role in ACD is the design of a deployable infrastructure that will allow a properly authorized defender to set up and initiate a defensive response to the threat. Implementation of the platform will be dependent upon collaboration and agreement with network owners. Our specific focus is defending DoD internal networks (for example, from the network boundary into and including the host) through the integration and automation of existing cyber-security solutions. The motivation behind the IAD's focus is derived from

- National Security Directive (NSD) 42 (The White House 1990) that "establishes initial objectives, policies, and an organizational structure to guide the conduct of activities to secure national security systems from exploitation; establishes a mechanism for policy development and dissemination; and assigns responsibilities for implementation",
- Executive Order (EO) 13587 Independent Assessment (The White House 2011) states "the Secretary of Defense and the Director, National Security Agency, shall jointly act as the Executive Agent for Safeguarding Classified Information on Computer Networks" and,
- the IAD's decades of experience defending DoD networks.

Attacks in the non-cyber domains require physical proximity and time to execute (for example, a bomb must be close to a target; a bullet must physically hit its target). Cyber is unique in the lack of need for physical proximity to execute an attack (that is, anyone with an Internet connection is a potential participant in this worldwide battle space) and in the vastly reduced time required to perpetrate an attack (for example, bits on a wire travel much more quickly than traditional troops or munitions). ACD addresses this vastly reduced time necessary for a successful attack by integrating many solutions to provide response actions in cyber-relevant time.

Cyber-relevant time is a purposely vague term that accommodates the needs of the battle space. If the battle space is a Central Processing Unit (CPU) and Random Access Memory (RAM), and the combatants are software applications vying for control, the cyber-relevant time is nanoseconds to microseconds. If the battle space is between two computers of close physical proximity, cyber-relevant time is milliseconds to seconds. For a battle space between two computers on opposite sides of the world communicating via satellite links, cyber-relevant time is seconds. With live operators and delays inherent in cognitive processing, key strokes, and mouse clicks, cyber-relevant time is seconds to minutes. The requirements for ACD increase as the adversary becomes smarter and quicker.

Cyber defense includes employing non-real-time big-data analytics to find trends in historical data repositories; likewise, cyber defense includes actuarial-like predictions of future events. The

ACD monitoring activity may provide data feeds to these analytics, and the ACD sense-making activity may take influence from these analytics in the form of decision support algorithms; however, these historical and future analytics are outside the scope of real-time processing and, therefore, outside the scope of ACD.

ACD as a Capability

A comprehensive ACD solution requires the integration of many tools. The complexities of ACD can never be entirely captured in a single tool. ACD functionality may occur within a single platform, but this is one example or one thread through ACD and not the entirety of ACD. Moreover, ACD functions may be geographically dispersed: sensing may occur in Hawaii; sense-making may occur in Washington, D.C.; decision-making may occur in U.S. Cyber Command (USCYBERCOM), and acting may occur in the European Command (EUCOM). The ACD design must accommodate a wide spectrum of such scenarios with performance occurring in cyber-relevant time. Therefore, the approach is to design ACD as a capability expressing desired results that may consist of an indeterminate number of tools that provide those results.

The primary beneficiary of ACD is the decision-maker. Decision-making is the act of selecting the best choice(s) among available options. Each decision-maker receives guidance from decision-drivers in the form of externally imposed authoritative mandates (legislation, regulation, directive, instruction, Executive Order), negotiated mandates (contracts, service-level agreements), or self-imposed mandates (internal policy, standards, procedures). Deriving ACD requirements includes decomposing each decision-driver into data elements necessary to make the decision. For example, operations standards for a particular mission may include requirements for specific values in a series of Windows registry entries. The ACD administrator decomposes this operations standard to identify all the parameters that represent compliance with that standard. The parameters represent requirements for decision-support, which is the information necessary for the decision-maker to decide.

The next step is to translate the decision-support parameters into data sources in the asset space and source data on each respective asset. Once the data needed and their location in the cyber-asset space are identified, appropriate sensors can be deployed to retrieve the data. Decision-makers provide requirements for the content necessary to collect (as just described) as well as the necessary frequency of data collection. The guiding principle to determine frequency is the question, 'How old is too old to make an effective decision?' For a compliance decision that occurs annually, collecting once a year is adequate. For critical operational decisions under emerging threat conditions, collecting once every X seconds or minutes may be more appropriate. This variety implies the need to dynamically configure sensors to accommodate changing mission needs. An increase in collection data volume and frequency will affect cyber environment performance. Because the cost of increased security can be decreased mission performance, there is a tradeoff between security and ease of use/performance.

ACD follows the principle of 'collect once and reuse many'. This means that data elements collected for ACD will be used for multiple purposes, including decision-support to areas outside of ACD. For example, the same data elements used to ensure operational security may also help to determine Federal Information Security Management Act of 2002 (FISMA) compliance. In some cases, the primary motivation to collect data may be for ACD, and these same data may be

useful in these other areas, such as certification and accreditation decisions, Command Cyber Readiness Inspection (CCRI) decisions, or FISMA compliance decisions. In other cases, the primary motivation to collect data may be for some other reason (for example, to validate operational service-level agreements), and these same data elements may be useful to ACD. Whatever the initial motivation to collect data from cyber assets, there is an even greater motivation to reuse that data in as many decision paths as it is applicable. This goal promotes smart and efficient workflow. Achieving this efficiency requires aware and intelligent management of integrated cyber defense to understand the needs and function of the individual parts in the context of the entire operation.

ACD Constituent Parts

ACD consists of six functional areas shown in **Figure 1:** sensing, sense-making, decision-making, acting, messaging and control, and ACD mission management. These logical functional areas are used to delineate role, fit, function, and dependencies. Sensing is ongoing observation with intent to provide awareness. Sensors are devices or people who make these observations and obtain a snapshot of current operational states. Sense-making uses analytics to provide understanding in a particular context (for example, mission, operational need, local security configuration). Each decision-maker will have a unique context within which to make decisions. ACD accommodates the automation of decision-making as well as the cognitive supplement of human decision-makers (that is, ACD may provide decision support).

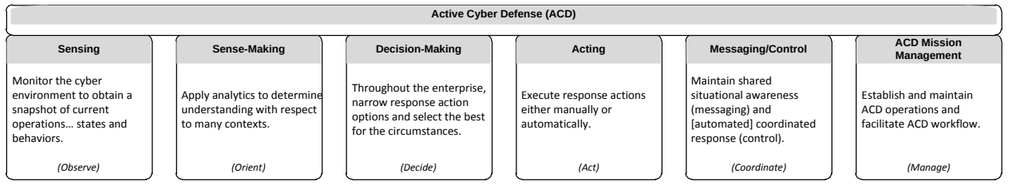

Figure 1: ACD Functional Areas

The objective of decision-making is to select among available response actions. ACD accommodates the ability to execute ACD-internal-automated-response actions as well as the ability to prompt external actors with action recommendations. The action decision ultimately resides with the cyber-asset owner. ACD does not impose *de facto* actions out of the control of the asset owner. While ACD calls out a decision-making function, decisions are made throughout the ACD workflow and outside of ACD using input from ACD. Decision-maker roles include, but are not limited to ACD administrators, cyber-asset owners, mission commanders, and security-operations personnel. Each has his/her own context and decision-drivers. There remains the challenge to establish real-time precedence and adjudication to resolve inevitable conflicts (for example, national policy requires X, but local operating needs require Y) to determine who wins and provide defensible justification.

Messaging and control is the heart of situational awareness and coordinated response actions. Key operational gaps in ACD are the lack of a common communication medium (for example, message fabric) to interconnect all ACD-related tools at speed and scale, the lack of a standard interface for tool connection to the common communications medium, and the lack of a standard

message set understandable and actionable by all connected tools. Upon successful realization of messaging and control, all ACD-related tools will have the ability to make each other aware of current activity (that is, to achieve shared situational awareness). Similarly, messaging and control will enable the tools to coordinate response actions that may include disseminating the same response action among similar assets (that is, a vulnerability mitigated in one is a vulnerability mitigated in all) or a more sophisticated combination of defense actions that hit multiple layers of network defense to preempt adversary attack and/or prepare the enterprise to weather an active attack.

ACD mission management provides ACD internal control of workflow where the scope of control is limited to the operating environment of that particular instance of ACD. There is no universal management of ACD. The ACD Reference Architecture (currently in draft) is to guide many instantiations of ACD—some of which will be standalone operations, and some of which will connect in varying degrees of coordinated operations. Participation in any semblance of a federated ACD operation is purely voluntary, and participants choose their level of participation. Participation is not imposed and certainly does not take place without the knowledge of the cyber asset owner(s).

To reiterate, ACD is not a single solution; it is a capability to provide context and interoperability among many solutions under the six functional areas. An integrated, cohesive ACD solution implies the use of many sensors, analytics, and displays to support many decision-makers. For example, ACD may accommodate any number of analytics from any number of perspectives. The type and focus of the analytic is dependent upon the needs of the decision-maker who will use the results of the analysis. ACD intends to accommodate what is available today (current tools) and what will be available tomorrow (future tools yet unknown). This leaves room for new, better, faster, and cheaper solutions across all functional areas.

ACD Operational Concepts

No single government entity will own ACD. ACD is a capability within cyber defense with the unique differentiator of providing situational awareness and response actions within cyber-relevant time. The only way to achieve this is to integrate many dozens of tools across the ACD functional areas. Sensing will include sensors and sensor subsystems (that is, sensor-management systems). ACD may have some native sensors (that is, sensors controlled by an ACD instance), but ACD will more likely interact with sensor-management systems that, in turn, directly control the sensors. One example of this is the Host-Based Security System (HBSS). ACD does not intend to directly touch any HBSS sensor; rather, ACD communicates with the ePolicy Orchestrator (ePO) server that, in turn, controls many HBSS sensors. A variety of sensors and sensor subsystems are necessary to monitor cyber assets (for example, Windows desktops and servers; UNIX; mainframes; network infrastructure, including routers and switches from many vendors; phone equipment; mobile equipment; industrial equipment, including Supervisory Control And Data Acquisition (SCADA)). No one sensor or sensor subsystem can watch them all.

Sense-making includes a wide spectrum of analytics that convert raw data into information for decision-support. Decision-makers are at every organizational level and include computer operators, system administrators, operations managers, leveraged security services, program

managers, investment managers, policy makers, and governance. NASCAR provides a useful analogy. In NASCAR, the car driver is the operator and is looking for details on speed, proximity of other cars, fuel level, tire pressure, laps-to-go, and position relative to the leader. The pit crew provides technical support in refueling, changing tires, and repairing the engine. The race track owner is concerned with track facilities, racetrack schedules, parking, and attendee safety. The racing commission is concerned with marketing, maintaining policy, coordinating all race schedules, ranking drivers, and providing overall governance. Each of these decision-makers is related under the sport of NASCAR. However, each has very distinct roles and distinct decision-making needs. Moreover, as NASCAR is not the only type of car racing, car racing is not the only type of sport, and cyber is not the only war-fighter activity. ACD will support a subset of these decision-makers. In fact, data collected by ACD may support more decision-makers outside of the scope of ACD.

An important distinction between contexts in the above decision-making analogy is that the racing commission really does not have a need to see all the speedometer readouts of every car actively racing on every track. Likewise, it makes no sense to display an average speed of every car actively racing on every track. There is no operational decision to be made from this information, and the racing commission should not reach out and step on the gas pedal or stomp on the brake of any particular car as that is the job of the driver. However, the racing commission is very much interested in collecting information on car performance, pit crew performance, and race-track results to make policy decisions that further the overall interest, performance, and safety of stakeholders. Of particular interest to the racing commission is profitability, which implies overall cost management and operational efficiency. While USG use of ACD is not concerned with profitability, ACD does contribute to cost management and operational efficiency via automation.

Upon receipt of decision support from sense-making, decision-making selects the best choice(s) among available options that ultimately leads to some action. Acting is the performance of a sequence of steps resulting from choices made in decision-making process. Actions may be manual or automated. A key principle in ACD design is for actions to be automatable and not inherently automatic. Automatable may be considered similar to a plane's automatic pilot. Under certain circumstances, automatic pilot is useful; however, the live pilot makes the choice to use it or not. ACD operates under the same principle. If an operator is comfortable with certain actions being automated, then that operator can flip the switch on and let ACD do its thing from beginning to end. If operating conditions are such that any automatic change to the operating environment is undesirable, then the operator will flip the switch off and queue ACD-recommended actions for review, approval, and execution by the live operator.

Messaging and control is the ability for ACD to provide shared situational awareness via standard communication methods and to provide coordinated response actions via standard control signaling with a standard message set. Shared situational awareness is informational only; that is, 'Here is a heads up on what we see' or 'Let us talk if this is of interest'. Coordinated response actions are technical devices informing other technical devices of recently performed or imminent actions and requesting/directing these other devices to take action as part of an overall coordinated response.

ACD mission management covers the establishment and maintenance of overall ACD, and it facilitates workflow through the functional areas. For example, if a sensor needs updating, ACD mission management handles the update; if a new analytic becomes available, ACD mission management inserts it into the appropriate area. Moreover, ACD facilitates workflow through sensing, sense-making, decision-making, acting, and messaging and control. By separating out the management functions, each functional area may focus on its particular role. This separation increases design and operating efficiency by isolating common functions under mission management versus duplicating them in each functional area.

ACD's Role in Broader Cyber Operations

ACD is a part of overall cyber defense that itself is but part of the broader cyber operations in support of mission execution. From a DoD perspective, the overarching concept is Computer Network Operations (CNO) that consist of Computer Network Attack (CNA), Computer Network Exploitation (CNE), and Computer Network Defense (CND). The current DoD Instruction 8530 (2001) addresses CND, and the latest draft revision renames CND to cyberspace defense. Any addition of a qualifier by nature reduces the scope (for example, the set of red cars is a subset of cars: all red cars are cars, but not all cars are red). Similarly, ACD is a subset of cyber defense: all ACD is cyber defense, but not all cyber defense is active.

The Operating Environment

From a DoD perspective, ACD exists within a federated operating environment in which most decisions with regard to the cyber assets are made by the respective asset owners. ACD may recommend and may facilitate automated responses that the asset owners agree to; however, ACD will not override asset-owner decisions. An important point is that ACD does not come prepackaged with foregone conclusions that will take over local network operations.

Operational Goals and Objectives

The cyber attacker's goal is to generate a desired effect brought about via three objectives of get in, stay in, and act. The cyber defender's goal is to minimize threat efficacy brought about via three objectives of keep out, throw out, and restrain. The desired effect may be tactical (for example, destroy a database server), or it may be strategic (for example, disrupt troop deployments that result from the database server output). ACD provides support to defend against both by first identifying and defending tactical targets and providing situational awareness to tip and cue mission-assurance-related activities.

General Cyber Attack Sequence

A generic attack sequence to get in is to enumerate the cyber environment, find vulnerabilities, gain access, escalate privileges, insert malware, and operate malware to the desired effect.

A generic set of stay-in activities are proliferate, avoid detection, and persist. Proliferate implies malware duplication with the intent of hedging attacker odds against detection of any single malware instance. Avoid detection implies hiding (for example, rootkits that insert hooks and modify operating system commands or common processes such as dynamic link libraries). Persistence is surviving through various conditions of rebooting, software patches, and other system modifications.

A generic set of act activities is designed to perform function and produce results. Perform function is the running of the exploit. The nature of exploits varies widely and includes every kind of malware that may attempt unauthorized disclosure of data, denial of service, or unauthorized modification of data. The production of results includes both a tactical result (for example, destruction of data) and a strategic result (that is, mission implication). ACD predominantly addresses the tactical result; however, parts of ACD work in complement with mission assurance to detect and respond to the adversary's intended strategic result.

ACD intends to monitor for the presence, state, and behavior of attacker attempts to get in, stay in, and act. For example, ACD sensing looks for enumeration behavior on the network (that is, activity that is mapping the network). Sensing looks for behavior that is attempting to identify vulnerabilities and monitor behavior as well as states that indicate unauthorized access, unauthorized privileged-user presence and activity, the presence of malware, and the activity of malware. Upon detection, the ACD workflow continues through sense-making, decision-making, and acting.

ACD Operational Example

An operational ACD capability might be a system of individual cyber-security solutions already deployed on a network (for example, Intrusion Detection System (IDS), Intrusion Prevention System (IPS), virus scanner, adaptive firewall) integrated together to produce an enterprise-spanning, holistic ACD capability. ACD may also be a set of unique tools integrated in a single platform to provide active defense against a specific threat vector in a specific portion of a network (for example, network or enclave boundary, host). Additionally, while mature solutions will ideally be vendor agnostic, based on common open standards and common command-and-control messaging, early instantiations of ACD will by necessity have proprietary elements.

As an early example of a platform-based ACD capability, the SHARKSEER solution is a collection of best-in-breed commercial products integrated into a single suite that provides active defense against zero-day attacks. This innovative capability includes a non-signature-based network sensor that identifies known malware in near real-time using government-enhanced, commercially developed signature and heuristic cloud technologies, and two behavior-based sensors where one sensor focuses on identifying real-time malicious human driven behaviors and the other on malware behavior. SHARKSEER is integrated on a state-of-the-art platform that provides high-speed, low-latency communication among the component parts and adds the ability to block malicious connections. This government enhanced integration of commercial products is an early instance of an operational ACD system designed by the IAD for defense of DoD networks and is capable of identifying and defeating rapidly evolving, previously unknown attacks that the individual products working alone cannot defeat.

To measure the fit of the SHARKSEER solution, sensing is first examined using the ACD functional framework. The SHARKSEER solution uses network flow sensors to route live network traffic through three streaming analytic capabilities. In sense-making, malicious-human-behavior analytics identify antecedent behaviors on incoming connections that are related to attempts to compromise the network. Real-time, signature-based analytics examine incoming traffic for indicators of known malware or files with bad reputation scores. Malware-behavior-based analytics examine incoming files and Uniform Resource Locators (URLs). In decision-

making, alerts from the sense-making analytics stimulate simple courses of action passed to acting. In acting, SHARKSEER really performs two basic functions: blocking and passing packets. The messaging and control function is provided by the hardware platform hosting the SHARKSEER solution.

SHARKSEER is intended for deployment at the network boundary (between internal and external networks) or at enclave boundaries (subdivisions of the internal network) and is designed to integrate with enterprise-wide holistic ACD solutions as other capabilities emerge. SHARKSEER is already integrated with enterprise email sensors and host-based systems via two deployed private clouds. Also part of the DoD Joint Information Environment Single Security Architecture and the Joint Regional Security Stack, SHARKSEER is slated for near-term deployment to provide defense for critical DoD networks.

Conclusion

The IAD's work on ACD is complete when ACD becomes a living part of DoD cyber operations. At that point, ACD will be in a maturity model that upon reaching the highest level of maturity will need to evolve the functional areas to sustain itself at that level. Part of this maturity considers that future requirements for ACD will evolve that cannot all be anticipated today. For this reason, the ACD architecture and systems engineering are capability-based and not tool-based. Capabilities are expressions of desired results, agnostic of the solutions that produce those results. Tools come and go with changing technology; capability needs are more enduring. The tools of ACD today will very likely be different ten years from now; however, the desired results from ACD should be close to the same. The IAD's ultimate vision for ACD is a capability that becomes a living part of how DoD operates and is conducive to emergent behavior so that the results of ACD as a whole are more than the behavior of the sum of its parts.

Way Forward

The way forward for ACD from a functional perspective involves dozens of parallel activities spanning all ACD functional areas. These activities occur predominantly in commercial development, some with and some without government sponsorship or explicit government requirements. A key activity to produce a cohesive ACD capability is integration. This implies the need for a common communication medium (for example, message fabric), standard interface, and standard message set. The goal is to adopt, adapt, or develop this common communications medium, standard interface, and standard message set for vendor use to create products that may become part of standardized ACD. Each individual product brings its unique value, and the whole of ACD becomes more than the sum of the parts as tool interoperability provides for the realization of cohesive and adaptive ACD operations.

Achieving the vision of a standardized ACD includes engaging the National Institute of Standards and Technology (NIST), commercial vendors, industry leaders in security concepts and technology research, and appropriate USG governance bodies across DoD, civil agencies, and the intelligence community.

ACD is a first glimpse at the broader capability of security automation, which is to maintain a state of being free from danger or threat within acceptable risk-tolerance boundaries with little or no human intervention. The benefits include workflow efficiencies, process coordination, priority

task execution, and intelligent resource allocations. The potential pitfalls include race conditions, gridlock, thrashing, and subverting parts of security automation for means other than their designed intentions. A clever adversary may turn poorly designed security automation into an attack tool that works against itself. These examples do not discourage security automation; rather, they raise awareness for careful security automation design, including the design of ACD.

References

Department of Defense 2001, *Department of Defense Directive 8530.01, Computer Network Defense*, Department of Defense, Washington, D.C., United States.

—2011, *United States Department of Defense Strategy for Operating in Cyberspace*, Department of Defense, Washington, D.C., United States.

The White House 2011, *Executive Order (EO) 13587, Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing of Classified Information*, The White House, Washington, D.C., United States.

—1990, *National Security Directive (NSD) 42, National Policy for the Security of National Security Telecommunications and Information Systems*, The White House, Washington, D.C., United States.

Verizon 2013, *2013 Data breach investigations report*, viewed 5 February 2014, <www.verizonenterprise.com/DBIR/2013>.

Securing the Cloud

D Parr

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *This paper will review cloud technology utilized to support the Intelligence Community and will specifically address the National Security Agency's research into vulnerabilities and risks related to cloud-based systems. Current implementation plans will be discussed for a multi-agency private cloud architecture that is under development. The paper will also review security challenges for a cloud architecture and will address specific technologies, such as data tagging, digital policy management, encryption, identity and access management, and auditing, along with intrusion detection and prevention.*

Keywords: *Cyber, Cloud, Cybersecurity, Encryption, Data Tagging, Intrusion Detection, Intrusion Prevention, NSA, CIA, ODNI, JIE, IC ITE, C2S, CNCI*

Introduction

The Cloud. Everyone is talking about it. Cloud computing is demonstrating many benefits to the commercial sector for savings and efficiencies resulting from economies of scale in the sharing of resources of power, space, and cooling. In addition to freeing up the corporate Information Technology (IT) organization from day-to-day operations of large computing systems, it is also credited with driving innovation and lowering costs. The U.S. government and more specifically national security organizations need to leverage this same technology to optimize computing power just as private industry is doing. The vast amount of data stored and utilized on a daily basis to fulfill national security missions lends itself well to efficiencies gained by cloud infrastructures and virtualization. The added requirement to collaborate across organizations to share national security information further promotes a shared computing environment with accessibility across agencies for those working specific problem sets.

The current trend towards cloud computing is an area of great interest to NSA and the rest of the Intelligence Community (IC). While U.S. government systems cannot always utilize commercial cloud services available to the general public, this same technology can be brought in house to leverage the computing power, resiliency, on-demand resources, and numerous other benefits of cloud technology. Gartner Group, a leading information technology research and advisory company, forecasts that cloud computing will become the bulk of new IT expenditures in 2016 (2013). The federal government is applying its technology budgets towards some of these same expenditures to take advantage of computing efficiencies. However, care must be taken to build these systems with assurance and trust. NSA's Information Assurance Directorate (IAD) is heavily involved in projects related to security for cloud architectures to meet the future

computing needs of the Intelligence Community. NSA is leveraging this technology for optimum advantage while providing confidence in data security.

While the benefits of cloud computing are evident, the security of these systems cannot be overlooked. NSA's IAD, chartered with protecting U.S. National Security Systems (NSS), is heavily involved in projects related to security for cloud architectures. It is important that security be considered in the early stages of development of cloud architectures. In a white paper regarding the securing of private clouds, Trend Micro, a Japanese global security company, cautioned that "The need for security must not be overlooked or 'bolted on' later during the transition to private cloud computing" (2014). Another company specializing in this field and a recognized leader in security solutions and provider of threat research, McAfee has released its *McAfee Labs 2014 threat predictions,* which reads, in part, "deployment of cloud-based corporate applications will create new attack surfaces that will be exploited by cybercriminals". Just as bank robbers go where the money is kept, so "Cybercriminal gangs of the 21st century will target cloud-based applications and data repositories because that's where the data is, or will be soon enough" (Intel 2014).

This paper will outline various NSA positions on cloud security and will review some of the technical focus areas. It will briefly discuss two major program initiatives, one led by the IC under the leadership of the Office of the Director of National Intelligence (ODNI) and the other by the Department of Defense (DoD). An overview of NSA's top vulnerabilities and mitigations for this growing technology area is included. The paper reviews specific techniques for protecting data and cloud platforms as outlined by NSA's IAD including data tagging, encryption, authentication, and other areas that are essential for protecting a cloud-based infrastructure. Finally, the threat situation is discussed with real examples of cyber hacking on data centers and cloud-centric computing environments.

ODNI Intelligence Community Information Technology Enterprise

To provide IT efficiencies with goals of improved collaboration and cost savings, ODNI is leading an initiative to bring cloud-based services to the top-secret national intelligence network. IC IT Enterprise (IC ITE) (pronounced 'eye sight') moves the IC from agency-centric computing to a common, cloud-based platform, allowing sharing of technology, information, and resources across the community. IC ITE will deliver more innovative and secure technologies to the desktop that will connect a large number of users across the IC.

IC ITE reached operational baseline in August 2013 with NSA providing the first portion of the utility, data, and storage cloud services combined with desktop services offered by other intelligence agencies. This news was recently announced to journalists and reported by *Federal Times* in an article titled 'Intel agencies roll out common IT services'. The article includes this assessment from Al Tarasiuk, Chief Information Officer for the ODNI and the person responsible for 'spearheading' the ODNI-led effort, '"We achieved a pretty significant milestone"' by reaching operational baseline (Johnson 2013).

The IC ITE cloud architecture is building access and redundancy beyond what was previously offered to the IC. Multiple agencies are participating and delivering services including desktop services, enterprise management, data and application hosting, utility, storage, and security and identity management.

A commercial cloud service called Commercial Cloud Services (C2S), a system in use by the Central Intelligence Agency (CIA), is supporting the utility and storage hosting portion of IC ITE. C2S is housed in a private data center on government premises. C2S features automated load balancing, the ability to provision new resources in just 900 seconds, usage and security reporting, and redundancy across multiple locations. It is available from any computer that connects to Joint Worldwide Intelligence Communications System (JWICS) and includes a multi-pronged security approach including

- Certification and authorization following current security regulations,
- Physical security with highly controlled data centers,
- Secure, restricted access and control of applications, and
- Data privacy including data encryption.

DoD Joint Information Environment

The Joint Information Environment (JIE) is a program managed by DoD's Defense Information Systems Agency (DISA) and supported by NSA. It is not specifically a private cloud environment such as the ODNI-led program, but instead a new architecture for a joint-information processing environment for the DoD secret-level network. JIE will leverage existing data centers, while consolidating and generating efficiencies of scale. JIE is comprised of a shared IT infrastructure, which may include cloud services, enterprise services, and a single security architecture to achieve full-spectrum superiority, improve mission effectiveness, increase security, and realize IT efficiencies. JIE is operated and managed by DoD command organizations using enforceable standards, specifications, and common tactics, techniques, and procedures.

The major goals for the JIE effort include

- Network optimization using common network standards,
- Single security architecture,
- Data center consolidation (core enterprise data center standards and consolidation),
- Identity and access management (attribution while supporting accessibility),
- Enterprise services (common capabilities across departments while still supporting mission specific applications), and
- Governance (consolidation of military services plans into optimized DoD plan).

NSA, through the IAD, serves as the Security Advisor to the JIE providing guidance on all cyber-security-related matters including security architecture development and coordination with U.S. Cyber Command (USCC) in their cyber-defense role, as well as serving as primary advisor and lead cyber-security consultant for the development of the JIE security architecture.

The ODNI IC ITE program and DoD's JIE are just two representative samples of the many initiatives underway within the national security community to leverage cloud technology. These systems along with other cloud-based systems under development need to consider security in their design and operations to provide the best possible protection of highly sensitive data. Data stored in these systems is utilized for major decision making by policymakers along with programs involving life and death situations. Data security is something that cannot be taken for granted.

The following sections will review some of the known vulnerabilities within cloud computing systems and possible mitigation techniques.

Key Security Topics for Cloud Computing

A secure, defensible computing environment requires a holistic approach to security that addresses organizational processes, policies, and technology, all with a focus on adequately protecting data and services. A cloud architecture has multiple layers and seams which introduce varying degrees of vulnerability. The diagram below describes NSA IAD's view of the major stacks of such an architecture and is itself followed by a discussion of the vulnerabilities and exploitations of those layers.

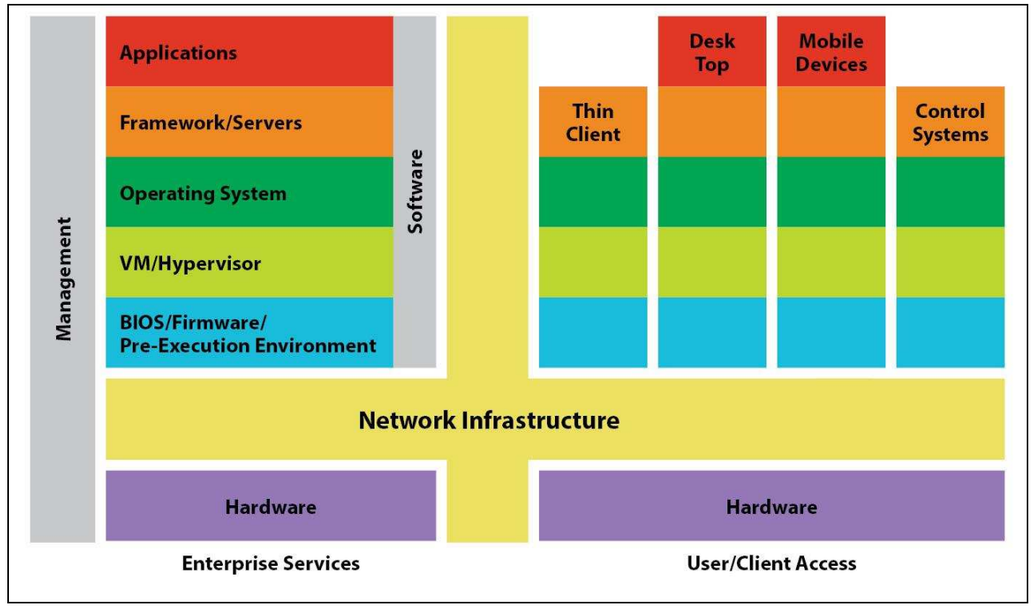

Figure 1: Cloud Environment Stack

Vulnerabilities/Exploitation

The stacks in **Figure 1** review the various surfaces of vulnerability. While many of the same vulnerabilities exist in non-virtual environments, issues such as the introduction of rogue devices become more risky in a cloud system. The following list reviews some of the more obvious vulnerabilities along with a few of the associated gaps in mitigation tools and techniques.

Hardware

- Introduction of rogue devices, implants
- Insider threat
- Supply-chain risk

Firmware/Basic Input/Output System (BIOS)

- Exploitations to gain persistence or disable devices
- Virtual Machine (VM) Manager/Hypervisor exploits plus infrequent patching of VMs
- Break out and break in attacks, multi-tenancy concerns

- Storage-management exploitation

Operating System

- Operating System (OS) image attacks
- Infrequent patching of OS
- OS exploits (no different from legacy computing issues): integrity and modifications to OS allow insertion of malware in OS

Frameworks/Servers

- Exploiting server misconfiguration
- Inadequate monitoring and human involvement
- Exploitation of poor identification and authentication services

Applications

- Poor data control; unprotected data formats, lack of encryption and access controls
- Exploitation of design flaws allowing intrusions such as Structured Query Language (SQL) injection
- Rogue applications (mobile, desktop, others)

Network Infrastructure

- Network resiliency; network segmentation controls and concepts are immature
- Public cloud puts the network on the Internet with data centers potentially located outside of the U.S.
- Router and other network device compromises; supply-chain risk, counterfeit devices, poor patching procedures.

These topics outlined by NSA augment similar ones described by Trend Micro in its report on *Virtualization and cloud computing; security threats to evolving data centers* (2014). The company writes that some of the security risks in a virtualized infrastructure include "communication blind spots, inter-VM attacks, and mixed trust level VMs. Instant-on gaps and resource contention are also important considerations". Later in the same report, a scenario is described in which an attacker could compromise one guest VM which then passes the infection to other guest VMs on the same host. According to the Trend Micro report, "hyperjacking" is an attack during which malware penetrates one VM and then attacks the hypervisor:

> When a guest VM attempts this attack, it is often called a 'guest VM escape' because the guest VM breaks out of, or escapes, its isolated environment and attacks the host hypervisor. Once compromised, a hypervisor can then attack other guest VMs on that host.

Technology gaps

Numerous gaps exist for addressing vulnerabilities in cloud-based systems. For instance, better tools are needed to test, monitor, and measure items such as bad BIOS and rogue hardware device insertions. Host forensic capabilities in VMs are limited and will prove challenging for analyzing cyber attacks on VMs. Academia is starting to research the challenges of forensics in a virtualized environment. Technology traces can prove to be difficult along with potential legal challenges when remnants of an incident require extending into other customer environments. A virtual environment could mean that a forensics search might involve data physically stored on servers located in multiple countries. Different laws and legal authorities might exist. The end result is that vital information may be overlooked or inaccessible to investigators.

Other challenges exist for analytical capabilities delivered through technology with required interpretation by humans. Both industry and the U.S. Government are constantly struggling to educate more cyber analysts whose job is to monitor, synthesize, and analyze logs of audited results. While more skilled personnel are needed, the vendor community needs to build more automated tools that streamline notification and detection of cyber events. Technology that can compile reportable events into a clear, reportable format is just beginning to be developed.

Insider threat

The U.S. Government knows better than anyone of the risks of insiders with system-administration privileges. The co-tenancy of data center computing into one large physical location poses the risk of insiders, with administrative privileges, gaining access to data beyond their area(s) of authority. Cyber consulting organizations are working to prove these vulnerabilities exist and to push industry to make modifications to prevent or monitor for malicious activity, whether from an insider or external source. In a *Network World* September 2013 article, it was reported that programming author and consultant Jeff Cogswell has identified and written about a cloud-based system hack in a piece titled 'The Windows flaw that cracks Amazon Web Services'. According to Butler, Cogswell set out to prove that "Windows data volumes…in public clouds such as Amazon Web Services can be copied and have their access credentials modified, allowing a hacker to glean insights into the data" (2013).

Cloud industry advocates, including John Howie, Chief Operating Officer (COO) of the Cloud Security Alliance, did not believe the threat was significant but did acknowledge that an insider could do damage. Butler characterized Howie's response by stating "To execute the vulnerability, the hacker must have access to that data volume in order to be able to copy and manipulate it" (2013). This kind of attack would require the hacker to have access to the file storage, which would likely only occur by an insider. This is just one small example of the damage that can be done by an insider. Due to co-tenancy and shared storage systems that are inherent in cloud systems, a much larger amount of data can be accessed by people with administrative privileges.

Supply-chain risk

Building of cloud computing systems requires acquisitions of many different technologies from a vast array of commercial companies and open source consortia. Once a system goes live, any one component that is compromised can pose a significant risk to the remainder of the system. As part of globalization of the commercial information and communications technology marketplace, increased opportunities exist for those individuals, nations, and companies intent on harming the U.S. by penetrating the supply chain to gain unauthorized access to data, alter data, or interrupt communications. Risk stemming from both the domestic and globalized supply chain must be managed in a strategic and comprehensive way over the entire lifecycle of products, systems, and services.

To counter this risk, NSA and the rest of the federal government is following White House guidance outlined in the Comprehensive National Cybersecurity Initiative (CNCI) (Committee on National Security Systems 2012). Initiative #11 of CNCI was developed to create a multi-pronged approach for global supply-chain risk management. Several agencies, including NSA, are participating in the first effort to improve threat information sharing and vendor threat analysis across the IC, DoD, and some civilian government agencies to better inform their

procurement decisions and enhance program officers' ability to manage risk stemming from those decisions.

Managing this risk will require a greater awareness of the threats, vulnerabilities, and consequences associated with acquisition decisions; the development and employment of tools and resources to technically and operationally mitigate risk across the lifecycle of products (from design through retirement); the development of new acquisition policies and practices that reflect the complex global marketplace; and partnership with industry to develop and adopt supply-chain and risk-management standards and best practices. This initiative will also enhance federal government skills, policies, and processes and provide departments and agencies with a robust toolset to better manage and mitigate supply-chain risk.

Mitigation and Security Management

In order to provide a more defensible cloud-computing environment, mitigations must be addressed in the same way that traditional, legacy security is addressed. Solid systems engineering, governance, and security practices are still mandatory to provide assurance in computing systems. Additional vulnerabilities due to virtualization require added focus, including the following items:

- Device integrity – In order to deter persistence in the cloud environment, IAD has promoted various measures such as application whitelisting, anti-exploitation measures (for example, buffer overflow protection), and attestation to validate the boot environment (see NSA IAD's *Application whitelisting;* NSA IAD's *IAD security configuration guides).*
- Virtualization – Cloud architectures leverage the benefits of virtualization providing the benefits of efficient use of resources, isolation of services, and the ability to provision standard images for clients. Cons of this technology include multi-tenancy and vulnerabilities in the current software that allow attackers to exploit isolation and images.
- Seams – All networks have seams, points of transition within systems of systems and between layers. Seams are a particular focus in cloud architectures due to the amount of data that could be exposed at the seam between clients and data centers. Addressing seams requires analysis and implementation of control of these critical points.
- Data aggregation – Data centers are a key component of cloud and other centralized network designs. These centers are generally a discoverable, attractive target for exploitation (NSA IAD *IAD security configuration guides).* The virtualization and seams focus areas are manifested at the data centers where they can expose large amounts of data.
- Authentication and authorization – The multi-tenancy brought about by cloud architectures places a critical burden on the identification and authentication, and authorization systems to ensure clients do not exceed their authorized access. This requirement includes the need for device authentication to detect and prohibit rogue devices and applications. NSA has long recognized this challenge and has responded with Public Key Infrastructure (PKI), Smart Data, Trusted Data Format (TDF), and other measures.
- Security management and configuration – Security management covers many areas from control of hardware and software, to patching, to incident monitoring, to having responsible personnel identified for performing this function. Security technologies, such

as authorization and filtering, can be implemented effectively or ineffectively. The ability to manage every network element, allow the enterprise to securely configure nodes (NSA IAD *IAD configuration guides*), control access policy, tailor responses, perform effective patching, and collect pertinent audit data is foundational to a defensible architecture.

- Monitoring and encryption – Today's reality is that networks will be exploited, which makes monitoring and remediation essential. The use of encryption to protect data impacts the ability to monitor. Holistic-systems security engineering should be practiced to examine networks to balance these two necessary technologies. The enterprise must identify key indicators of compromise in cloud architecture and instrument the network to gain visibility. The management infrastructure must automate the consolidation and ingestion of this data so that innovation can focus on analytics to flag suspect behavior.
- Resiliency/active risk management – A defensible architecture needs the capability to respond to breach indicators by ramping up forensics, limiting or shutting off access, or executing other tuned responses within the enterprise as informed by both commercial data feeds and intelligence sources.

NSA IAD Contributions to Cloud Security Technology

NSA is embracing various technologies to make cloud computing for the DoD and IC private clouds more secure. While there are numerous security features in use, under development, and still in the research stages, here is a list of some of more mature technologies:

- User authentication, identity and access management;
- Access control to data elements;
- Digital policy management;
- Encryption;
- Auditing; and
- Monitoring and intrusion detection/intrusion prevention.

User authentication, identity and access management

Certain security services are considered foundational for a secure cloud infrastructure housing national security information. First, identity management creates and administers 'identities' that uniquely and unambiguously distinguish people and machines end-to-end across the enterprise. Credential management creates and administers 'credentials' based primarily on public key infrastructure (PKI), digital certificates to permit people and machines to prove their identities to access data and resources.

Attributes management manages and shares descriptions (attributes) about people and machines that are required to adjudicate accesses and validate attributes against their identity. It ensures only authorized individuals gain access to protected resources. Providing unified user attribution across the DoD and the IC is a challenge for such a broad community of interconnected users. To assist in this mission, the Unified Authorization and Attribute Service (UAAS), a joint DoD and IC effort, is providing the secure dynamic exchange of user attributes between participating agencies to support dynamic authorization decisions for IC and DoD users.

Another critical service is authorization management which adjudicates access requests to protected resources based on identities and attributes. More information about access control to data and policy management is part of the next section.

Access control to data elements

As cloud technology evolved within the IC, NSA quickly saw the importance of tagging data to control access to specific types of national security information. Some data, such as signals intelligence, requires specific handling and access controls based on legal authorizations beyond just a security clearance requirement. Monitoring for compliance and reporting must be rigorously maintained to adhere to legal mandates describing access to specific critical national security data.

Smart data tagging evolved to support this need to tag data and provide access control. These mechanisms have the requirement for tracking data movement, processing, and usage of the information. The owner of the data has the responsibility for identifying who meets required 'need-to-know' criteria. NSA's IAD created an open-source data-tagging library that is based on building a wrapper around data and its associated metadata. The format is known as the TDF and supports any data format. This format has been selected by ODNI as the IC standard. NSA intends to submit this standard to the open-source community as an industry standard. Here are some of the features of this data-tagging format.

- Secure container for data and metadata built with an open, self-describing format
  - Using Extensible Markup Language (XML) core specification, built upon and extending open standards
  - Specification extensible to Efficient XML Interchange (EXI), binary, and other encodings
- Integrity/authenticity (tamper-proof cryptographic binding)
  - Cryptographic protection of data, metadata, and relationships
  - Granular integrity verification
  - Signed assertions to support authenticity verification
- Encryption to support confidentiality
- Classification-marking support
  - Portion-mark, rollup, and tear-line capabilities to support sharing
  - Validation of tags against enterprise and mission business rules and policies to support accurate, consistent tagging
- Access Control
  - In-line with the IC ITE selected models
  - In-line with DoD-JIE's Enterprise Dynamic Access Control.

Figure 2 (below) depicts the Trusted Data Object format.

Digital policy management

NSA has developed an internal tool for automated building, managing, and approving of digital-access-control rules and policies based on data and user characteristics. This system will be directly utilized on the IC ITE program to enable greater integration and information sharing, while still safeguarding information through automated policy generation. The policies enforced by this system enforce attribute-based access control using a natural language interface that is easily understood by data stewards, mission managers, lawyers, and policy-compliance officers. The policies are stored in a central repository for use in protection of data and resources. The main purpose of this tool is to make sure that data is accessible by the widest audience with the correct clearances and 'need-to-know' requirements.

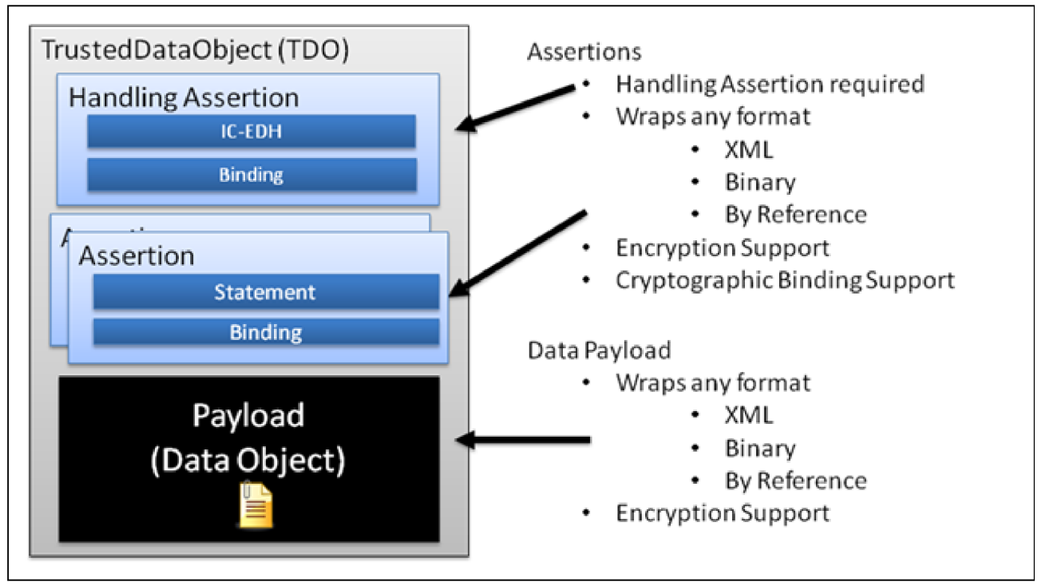

Figure 2: NSA Trusted Data Object Description

Encryption

As described earlier in this document, cloud-based computing adds numerous benefits while posing new vulnerability surfaces. Encryption of data is one of the methods to add layers of security to cloud-based content, whether it is a public cloud or private cloud system. In recent years, NSA has chosen to move to a layered commercial model for meeting national security encryption requirements. They have initiated the Commercial Solutions for Classified (CSfC) program that is delivering capability packages for industry to provide commercial solutions matching our encryption needs. Vendors who wish to have their products eligible as CSfC components of a composed, layered IA solution must build their products in accordance with the applicable U.S. Government Protection Profiles and submit their products using the National Information Assurance Partnership (NIAP) Process. More information about CSfC is provided in a separate paper in this issue of the *Journal of Information Warfare.*

One of the upcoming protection profiles to be released in 2014 is Data at Rest (DAR) security. DAR is the term used to describe computer data in storage, such that it is not traversing a network or temporarily residing in computer memory. DAR includes hard drives, removable media, and storage area networks. The CSfC approach for DAR is to protect any classified data stored in non-volatile memory with a solution that meets the following criteria:

1. Two layers of encryption using algorithms and key lengths as specified in the Committee on National Security Systems (CNSS) Policy No. 15.
2. Each layer of encryption provided by a commercial product that has been evaluated under the NIAP against a protection profile. Protection profiles for full-disk encryption and mobile devices have already been published. Protection profiles on file encryption and self-encrypting drives are expected to be published in 2014.

3. Solution contains cryptographically independent layers of protection, such that failure or compromise of two independent mechanisms is necessary to expose data.
4. Compliant with the requirements in the CSfC DAR capability package, version 1, expected to be published in 2014.

Auditing

The Enterprise Audit Management (EAM) program is a joint effort of NSA, USCC, and DISA to improve system-event detection on the DoD networks. The EAM system also includes inputs from the National Institute of Standards and Technology (NIST) and other DoD components.

EAM seeks to

- continuously detect and report enterprise events, maintaining situational awareness and responsiveness to threats in real-time;
- provide network speed alerts of detected anomalies;
- counter local and remote adversarial threats; and
- enable Active Cyber Defense (ACD) operations.

The EAM Program defines gaps in audit capabilities including monitoring and collection, analytics, alerting and reporting, storage security, and forensics. The program identifies options to close those gaps. It investigates the role and capabilities of audit in the current environment and in the future enterprise, including virtual, cloud, and mobile environments. As digital networks have grown larger, more complex, and more interconnected, identification of threat activities in time to respond with active defensive tactics has become more difficult. This program identifies tools and configurations to collect data on threat activity and analytics that can process this data and rapidly identify anomalies within the data. It outlines alert and presentation tools that quickly present notice of anomalies to administrators and analysts. The program also presents recommendations for audit policies, plans, procedures, and guidance in new environments. It is designed to address future cloud computing audit requirements.

Monitoring, intrusion detection and prevention

The DoD published the Strategy for Operating in Cyberspace in July 2011 with a definition of ACD as the DoD's "synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities" to defend the department's information networks (DoD 2011). NSA is employing commercial, open-source, and government-developed tools to address the need for automated cyber defense and incorporating them into monitoring, intrusion-detection, and prevention systems. These new services will be incorporated into various NSA, DoD, and IC cloud systems. Analytic tools and techniques are being incorporated into the services to provide automated alerting, reporting, and decision-making mechanisms to notify system administrators of malicious activity and prevent damage.

Cyber-Attack Threat

Threats to cloud-based computing are real. Numerous accounts have been documented of attacks to cloud-based systems. While vendors and government agencies such as NSA are working to quickly provide mitigation of the security vulnerabilities, hackers are only beginning to explore attack surfaces.

Amazon Web Services, a leader in cloud-computing services, was used by hackers in 2011 for an attack against Sony Corporation. According to a Bloomberg news article, "Hackers using an alias signed up to rent a server through Amazon's EC2 service and launched the attack from there". The incursion compromised the personal accounts of more than 100 million Sony customers and appeared to be very professional and highly sophisticated according to Sony. Jeff Bezos, CEO of Amazon, was quoted as saying, '"Data security is one of these great dynamic situations where the bad guys get better, and the good guys have to keep getting better too—it's not a static situation"' (Alpeyev, Galante & Yasu 2011).

A zero-day virtualization vulnerability was used in 2009 against a large Internet service provider based in the UK, according to a UK publication called *The Register*. Attackers were able to penetrate UK-based Vaservs servers by exploiting a critical vulnerability in HyperVM, a virtualization application made by a company called LXLabs. According to the report, data for as many as 100,000 websites was destroyed by attackers (Goodin 2009).

From a national security perspective, the U.S.-China Economic Review Commission warned in November 2013 that China's cyber-threat could come from the cloud. The report said, "a plan to link China's largest carrier-neutral internet data center services provider, 21Viant, to Microsoft data center in other countries suggest the Chinese government could eventually gain access to data centers outside China" (Chabrow 2013). With the growth of the global technology landscape, international interconnectivity will pose new security threats.

McAfee Labs summarized the extent of the problem in their *Forecast growth in mobile ransomware and security-aware attacks in 2014*:

> Deployment of cloud-based corporate applications will create new attack surfaces that will be exploited by cybercriminals. Cybercriminals will look for more ways to exploit the ubiquitous hypervisors found in all data centers, the multi-tenant access and communications infrastructure implicit in cloud services, and management infrastructure used to provision and monitor large-scale cloud services. (Intel 2014)

Conclusion

Cloud technology is evolving at a very rapid pace with new capabilities available every day. The U.S. government, like private industry, is leveraging the dramatic computing power of virtual systems for efficiencies and collaboration benefits. The vast amount of data stored and utilized on a daily basis to fulfill national security missions lends itself well to efficiencies gained by cloud infrastructures and virtualization. National security agencies are moving to cloud-based systems and challenging themselves to provide security for this sensitive data. NSA's IAD is working within the agency and with other government agencies to provide monitoring, mitigation, and awareness of cloud vulnerabilities. This paper outlines some of the techniques under development as well as new ones that are still maturing.

As the section on threats to existing cloud-based systems points out, hackers are only beginning to learn how to penetrate the new and exposed surfaces of this technology. As more data is moved to cloud systems, it is possible that new attempts to access this data will occur. This should not be a reason to avoid cloud computing as there will always be risks with physical and virtual computing systems. Great care needs to be taken as systems are designed and built. Jeff

Bezo's assessment bears repeating, '"bad guys get better, and the good guys have to keep getting better too'" (Alpeyev, Galante & Yasu 2011). As part of its IAD mission, the NSA will continue to provide expertise for protection of U.S. National Security Systems whether the data is stored in traditional physical computing systems or cloud-based virtual systems.

References

Alpeyev, P, Galante, J & Yasu, M 2011, 'Amazon.com server said to have been used in Sony attack', *Bloomberg News*, 14 May, viewed 15 January 2014, <http://www.bloomberg.com/news/2011-05-13/sony-network-said-to-have-been-invaded-by-hackers-using-amazon-com-server.html>.

Butler, B 2013, 'Programmer exploits Windows vulnerability in cloud-based services', *Network World*, 12 September, viewed 15 January 2014, <http://www.networkworld.com/news/2013/091213-windows-vulnerability-273793.html>.

Chabrow, Eric 2013, 'Shaming China to stop hacks doesn't work', *Data Breach Today*, 21 November, viewed 11 February 2014, <http://www.databreachtoday.com/shaming-doesnt-stop-chinese-hacks-a-6243>.

Committee on National Security Systems 2012, 'Use of public standards for the secure sharing of information among NSS', 1 October, viewed 31 January 2014, available through <www.cnss.gov/CNSS/issuances/Policies.cfm.>

Department of Defense 2011, *Department of Defense strategy for operating in cyberspace*, 14 July, viewed 28 January 2014, <www.defense.gov/news/d20110714cyber.pdf>.

Gartner Group 2013, *Gartner says cloud computing will become the bulk of new IT spend by 2016*, Gartner, 24 October, viewed 16 January 2014, <http://www.gartner.com/newsroom/id/2613015>.

Goodin, D 2009, 'Webhost hack wipes out data for 100,000 sites', *The Register*, 8 June, viewed 10 January 2014, <http://www.theregister.co.uk/2009/06/08/webhost_attack/>.

Intel Corporation 2014, *McAfee Labs 2014 threat predictions*, viewed 10 January 2014, <http://www.mcafee.com/us/resources/reports/rp-threats-predictions-2014.pdf>, p. 5.

Johnson, NB 2013, 'Intel agencies roll out common IT services', *Federal Times*, 9 September, viewed 13 January 2014, <http://www.federaltimes.com/article/20130909/IT03/309090015/Intel-agencies-roll-out-common-services>.

National Security Agency, Information Assurance Directorate n.d., *Application whitelisting*, Information Assurance Directorate, National Security Agency, viewed 28 January 2014, <www.nsa.gov/ia/mitigation_guidance>.

— *Information Assurance Directorate configuration guides*, Information Assurance Directorate, National Security Agency, viewed 28 January 2014, <www.nsa.gov/ia/mitigation_guidance/security_configuration_guides>.

Trend Micro 2014, *Adapting security to private cloud computing*, White Paper, Trend Micro, 1 January, viewed 7 February 2014, available from <http://www.gartner.com/technology/topics/cloud-computing.jsp>.

Trend Micro n.d., "*Virtualization and cloud computing: security threats to evolving data centers"*, viewed 7 February 2014, < http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/reports/rpt_security-threats-to-datacenters.pdf>, p. 3.

How IAD Leverages Big Data for Anomaly and Malware Detection (v10.2)

S Roddy

Information Fusion and Analysis Office
Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *Malware is growing increasingly sophisticated. Threats are becoming more targeted and moving to places where existing defenses have limited visibility. Proactively addressing these threats means leveraging insights gained from Big Data and the fusion of multiple sources of information. Operational Fusion and Analysis, OFA, an organization within the National Security Agency's Information Assurance Directorate utilizes Big Data to provide battlespace awareness and critical intelligence on the attack lifecycles of intrusions to decision makers and network defenders. This is accomplished by performing qualitative and quantitative analysis, summarization, fusion, and trending of data across multiple networks, customers, and domains. The more insight the OFA gains into a network or series of networks, the more easily abnormal activity can be identified.*

Keywords*:* *Big Data, Data Analysis, Malware*

What is Big Data?

Big Data is an emerging concept in the field of Information Assurance (IA). While marketing and search engine communities have been leveraging Big Data for years, these innovations are just recently being leveraged for network defense. From a National Security Agency (NSA) Information Assurance Directorate (IAD) perspective, Big Data comprises every bit of data and information that is relevant to protecting U.S. government critical networks. This includes everything from packet captures (PCAP) and log files to reporting and finished intelligence. By leveraging this vast sea of data, the NSA will not only counter current network threats, but also predict those that are emerging.

How Does IAD Use Big Data?

Identifying, characterizing, and trending Information Networks requires collecting and analyzing data such as host logs, network flow data, PCAP, vulnerability data, configurations, system and application data, and event data. The amount of relevant data, even for a single network, can be staggering, often several terabytes' worth for even the most basic snapshot.

NSA IAD's Operational Fusion and Analysis (OFA) division is charged with understanding and characterizing information systems. More specifically, OFA strives to use Big Data to characterize and report on networks' posture and vulnerabilities by understanding dataflow, structures, and usage. To meet this mission, OFA attempts to make complete use of the data and information at hand as efficiently as possible.

The three Vs

While characterizing Big Data as the three 'Vs'—volume, velocity, and variety—is not novel, it helps OFA understand Big Data and how to make sense of it.

Volume: Since the 1980s the amount of data has grown exponentially. Buried amidst this massive store of data are the important morsels of information that OFA needs to protect critical networks. The challenge for OFA, and the community, is that there is no sign of data volume leveling out, let alone decreasing.

Velocity: Velocity refers to the speed at which Big Data moves and needs to be analyzed. OFA's challenge is identifying the germane bits and making sense of them in military time to enable commanders to act. This is often one of the primary criteria used to differentiate Large Data from Big Data. Rasmus Wegener has provided a particularly useful analogy to illustrate the difference:

> "This is almost always a knockout criterion. When you walk through the airport and they take pictures of everybody in the security line to match every face through facial recognition, they have to do that almost in real-time. That becomes a big data problem. If I am a bank and looking at a vast number of credit scores and histories, and I don't need to provide an answer in five seconds but can do it next day, then that is not a big data problem". (Groenfeldt 2012)

Network threats arrive at the speed of light, and it is OFA's mission to enable Big Data sense-making to counter those threats either in real-time or ideally before an attack even occurs (predictive).

Variety: Variety means complexity. OFA must be able to parse the variety of data formats contained in Big Data. Some challenges include the ability to seamlessly analyze and fuse information contained in different formats. Handling the variety of formats must also be scalable as new formats emerge. Technology will certainly play a crucial role in keeping up with these formats and maintaining the capability to translate them into usable forms.

OFA's analysts work with numerous types of unstructured and multi-structured data from various sources. Unstructured data is unorganized or not easily processed by traditional databases or data models. Metadata, reports, and data files are examples of unstructured data. Multi-structured data takes Big Data to the next level; it consists of a variety of data formats and types created between human and machine interactions, such as web application logs:

> A great example is web log data, which includes a combination of text and visual images along with structured data like form or transactional information. A digital disruption transforms communication and interaction channels—and as marketers enhance the customer experience across devices, web properties, face-to-face interactions and social platforms—multi-structured data will continue to evolve. (Arthur 2013)

A fourth 'V,' Value, is sometimes included with the original three and refers not to the data itself, but to the information and knowledge that can be gleaned from the data.

Strategic focus of Big-Data analysis

In order to better align our Big-Data efforts, OFA developed four strategic focus areas. These areas serve as a guide for identifying information contained in Big Data that is critical for network defense and Information Assurance.

Structure: An awareness of how a network is engineered. It provides analysts with knowledge of host configurations, infrastructure configurations, and physical and logical mappings.

Usage: Simply put, what is happening on a network—good *and* bad. Usage analysis entails network activity characterization, log analysis, and user behavioral analysis.

Dataflow: Pertains to how data moves throughout a network including where the data goes and how much of it is moving. Dataflow analysis helps to determine exfiltration paths and allows analysts to develop related trends.

Posture: Determining Posture comprises comparative analysis, measuring vulnerability-mitigation effectiveness, and discovering exploitation trends uncovered by the other focus areas (Structure, Usage, Dataflow). While this is often the most interesting focus area, it is also the most difficult to determine and tends to be subjective.

Analysis methodology

In addition to utilizing the Strategic Focus Areas as guides, the analysis process illustrated in **Figure 1** below typically implements four primary techniques when working to glean critical information from Big Data:

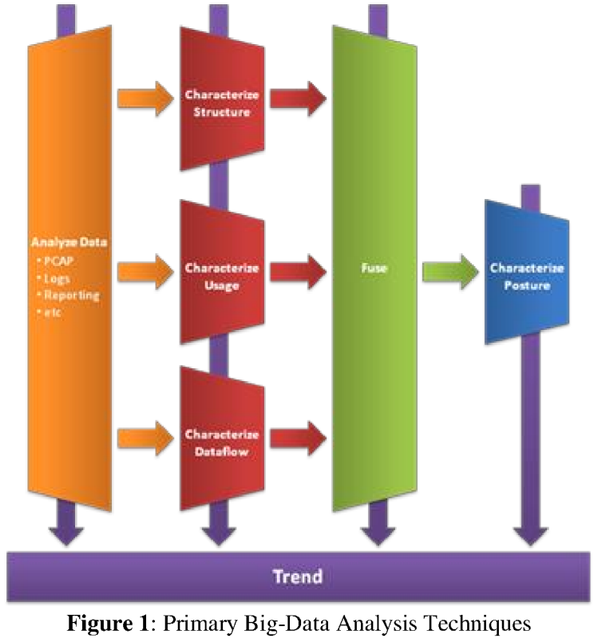

Figure 1: Primary Big-Data Analysis Techniques

Analyze: Analysis entails performing a deep-dive into the data at hand. Generally speaking, analysis will generate additional data or information as an output. An example of analysis could be identifying and extracting attributes (for instance, metadata) from executables.

Characterize: Characterization entails making some judgment as to what the data is stating. For example, this technique might take the executable metadata generated during analysis and categorize each file by its level of maliciousness.

Fuse: Fusion involves integrating multi-source data and then extracting key information. For example, fusion takes the metadata and categorizations generated during analysis and characterization, and integrates them with network activity to develop a more complete picture.

Trend: Trending involves extracting relevant information based on some change factor. Change factor examples include time, size, counts, etc. (For the malware example illustrated in **Figure 2** below, this could include trending infection rates over time.)

Big-Data Toolbox

OFA analysts have diverse backgrounds and possess a variety of skill sets; thus, an OFA analyst's tool box is a mosaic of software from myriad disciplines. Network analysis tools[1] are joined by forensics tools[2], productivity tools such as Visio, and a variety of NSA-developed tools. OFA analysts also use cloud-based tools similar to Pig to analyze large datasets, as well as virtual machines and personalized automated scripts. While some specialized tools are used, general-purpose tools tend to be preferred, especially for deep analysis, research, and knowledge-discovery purposes.

A mixology of analysts

The aggregate skill set leveraged by OFA is both broad and deep. Computer scientists of all experience levels are joined by intelligence analysts with 25-plus years of experience and recently graduated new hires in a comfortable environment in which senior and junior analysts work together to brainstorm and test new ideas. OFA analysts not only tackle Big Data sets with skills in network forensics, metadata analysis, and scripting, but they also strike with deep web research and traffic analysis to formulate the most complete story possible. OFA analysts with extensive experience in writing and graphic design ensure that any intelligence gleaned from Big Data is communicated to customers in a clear and concise form.

Coaxing actionable network intelligence from a mass of data requires a multi-disciplined work force. Stephen Sims (2012), Senior SANS Instructor, acknowledged the need for a combined quantitative-qualitative approach in a SANS Leadership Laboratory post:

> With information security, basing a final risk rating simply on numbers does not often result in the best analysis. Combining multiple elements gets us much closer to an accurate understanding of our threat level.

Positioned for a combined approach, OFA has assembled a team of analysts with experience ranging from packet analysis to geo-political target assessment. These analysts use their

[1] Some commonly used tools may include Wireshark, Splunk, SiLK, and Snort.
[2] Forensics tools may include EnCase.

understanding of cyber security and computer network defense as the common link for collaboration.

OFA Products Incorporating Big Data

OFA analysts have come up with unique customer-tailored analysis solutions that encompass both quantitative and qualitative methods. Two examples of their work include Cyber Battlespace Characterizations and a line of Malware Reporting delivered in placemat form.

Cyber Battlespace Characterizations (CBCs)

Cyber Battlespace Characterizations (CBCs) are one product that evolved from OFA's Big-Data analysis effort. These documents support myriad customers and offer insight into everything OFA knows about a particular network. This product is OFA's version of the reconnaissance work a hacker would do—except it is likely to be more complete and used for defensive purposes. Large volumes of data are sifted through by first accessing relevant sources of Big Data, chunked into categories, and analyzed and fused to extract meaning. The resulting 'story' or value is reported to OFA customers.

Malware placemats

Malware-placemat reporting is an example of OFA fusing information from qualitative and quantitative analysis. One example extracted and sanitized from an OFA report includes a fairly simple chart (**Figure 2**) using quantitative analysis:

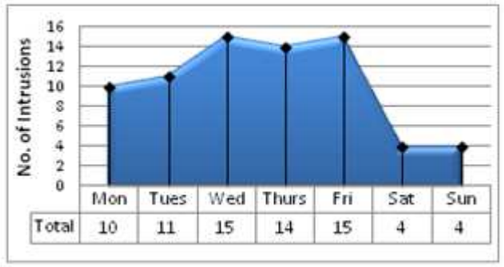

Figure 2: Sample Malware Infection Rate

The chart shows the infection rate of a particular piece of malware on a particular network. By pulling in a large amount of data from a wide-range of sources in a more complicated use of quantitative analysis, the analyst was also able to create a chart showing global activity levels for the malware. The global chart documented a decrease in infections, as well as which variant of the malware accounted for the majority of the activity. Qualitative analysis added further meaning by suggesting why the infection rate had dropped and proposed mitigation actions.

Big Data's Future in OFA and IAD

Since its inception, OFA has had a dynamic relationship with Big Data. Though vital to the mission, the huge quantity of data is expensive to process from both a manpower and information-technology perspective. Future success depends on the development of more refined sensing technology and backend analytics. Once in place, these innovations may permit OFA to leverage Big Data to predict and counter threats before or as they emerge.

References

Arthur, L 2013, 'What is Big Data?', *Forbes Magazine*, viewed 9 January 2014, <www.forbes.com/sites/lisaarthur/2013/08/15/what-is-big-data/>.

Groenfeldt, T 2012, "'Big Data" often isn't really, says Bain', *Forbes Magazine*, viewed 8 January 2014, <www.forbes.com/sites/tomgroenfeldt/2012/01/04/big-data-often-isnt-really-says-bain>.

Sims, S 2012, *Qualitative vs. quantitative risk assessments*, SANS Institute Leadership Laboratory, viewed 8 January 2014, <www.sans.edu/research/leadership-laboratory/article/risk-assessment>.

Stubbs, E 2012, *The value of high-performance analytics*, blog, 25 April 2012, SAS, viewed 9 January 2014, <www.blogs.sas.com/content/anz/2012/04/25/the-value-of-high-performance-analytics/>.

Tool References

EnCase, <www.guidancesoftware.com>.

Snort, <www.snort.org>.

Splunk, <www.splunk.com>.

System for Internet-Level Knowledge (SiLK), <http://tools.netsa.cert.org/silk>.

Wireshark, <www.wireshark.org>.

Taleb, NN 2007, *The black swan: the impact of the highly improbable*, Penguin Books, New York, NY, USA.

Outmaneuvering Cyber Adversaries Using Commercial Technologies

J Watkins

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *NSA characterizes assurance as having confidence that United States Government (USG) missions and networks can withstand attacks from determined adversaries. By Presidential Directive Order, the Director of the National Security Agency (DIRNSA) is the authority for National Security Systems (NSS), and this authority is delegated to the Information Assurance Directorate (IAD) to determine standards and policies for the nation's most critical security systems.*

Keywords*:* *IA, NSD-42, NSS, CSfC, GOTS, Open (non-proprietary) Standards, Designed-In Assurance Features, Diversity, Protection Profiles, Technology Communities, NIAP, CSfC Components List, Capability Packages, Risk Assessments, Approving Officials, CSfC Solution, Registration, NSA Mobility Program, Mobility Innovation Center, FISHBOWL Pilot, Trusted CSfC Integrators*

Introduction

Information Assurance (IA) has always been critical to United States Government (USG) customers, including the Department of Defense (DoD), members of the Intelligence Community, the military services, and other federal agencies. In today's increasingly interconnected world, Chief Executive Officers, Chief Information Officers (CIOs), and ordinary citizens are equally concerned about protecting their information, networks, and infrastructure. The IA mission has resided at the National Security Agency (NSA) for more than 60 years, and NSA carries out its critical role under National Security Directive-42 (NSD-42) with the Director of the NSA serving as the National Manager to safeguard National Security Systems (NSS). The NSA Director delegates this responsibility to the IA Directorate (IAD) to determine standards and policies for our nation's most critical security systems, and IAD protects NSS by hardening and defending systems that handle classified and other critical information. By adopting open standards, commercial industry builds products that are leveraged by banking, energy, transportation, and other critical infrastructures to enable their systems to operate securely. In this way, the success of the IAD's mission affects everyone who uses commercial IA products—both in the U.S. and around the world (*Confidence in cyberspace* 2013).

Confidence in Cyberspace

The IAD characterizes assurance as having confidence in the security properties of critical systems and networks even when they are under attack by an adversary. Today as new technologies are introduced at a rapid pace, the IAD remains committed to the mission of protecting information. As part of NSA's commercial strategy, the IAD's objective is to

determine how to design solutions for national security customers using commercial technology. The process known as Commercial Solutions for Classified (CSfC) is how NSA is executing its commercial strategy to influence and take maximum advantage of the benefits offered by commercial technology. This is achieved by layering commercial products to protect classified information, thereby enabling customers to communicate securely based on open, commercial standards (*Confidence in cyberspace* 2013).

A Cultural Change

CSfC is important today because users are demanding it. Given the current pace of technology, the IAD cannot keep up with the customer demand signals with only proprietary Government Off-The-Shelf (GOTS) IA solutions. Customers are impatient with government delivery schedules and increasingly frustrated by not being able to leverage current technology (*CSfC overview* 2014).

Given the popularity of tablets and smart phones, NSA's IA Director Debora Plunkett acknowledges customers are able to use those devices at home but '"come to their work spaces and it's like ... they can't do a whole lot. And we want to be able to change that"'. She said a '"cultural change"' is needed to meet government users' needs. '"Cell phones today are obsolete in ... 12 to 18 months"' (American Forces Press Service 2012).

'"We've got to be able to operate in that same cycle as we're looking at putting smart devices in the hands of government users"', Mrs. Plunkett said. '"We've got to be able to move quickly enough such that we can also be able to evaluate those new devices and put them in the hands of users in enough time while those devices are not obsolete"' (American Forces Press Service 2012).

Traditional GOTS development and evaluation processes just cannot keep pace with today's demand signals. '"[This] became the poster child, instead, for what we don't want to do"', according to Mrs. Plunkett. '"And that is have a very, very long cycle of development, five years to deliver—millions and millions of dollars to deliver—and by the time it comes to market it's been overtaken by technology"' (*Confidence in cyberspace* 2013). CSfC gives users the opportunity to leverage current technology while it is still current because industry can deliver products to market more quickly than the government (American Forces Press Service 2012).

Proper Configurations Are Key

Commercial solutions are sufficient for protecting classified information if they are used in a solution that is properly composed. This is the key to the IAD's success because the IAD must define the requirements and design the security architectures into which commercial technologies are composed (*CSfC* 2013).

Government customers frequently use commercial products to perform their mission, but they do not always have the optimal security settings configured. The IAD is committed to increasing functionality and usability through the appropriate use of commercial products. Although NSA's strategy for protecting classified information continues to employ both commercially-based and traditional GOTS IA solutions, the IAD looks first to commercial technology and solutions in helping customers meet their needs for protecting classified information (*CSfC* 2013).

Pros and Cons

Most commercial products are not built specifically with security in mind. Providing a rich set of features and a positive user experience are usually placed above security on most vendors' requirements lists. Commercial industry is very good at getting products to market quickly; however, those commercial development processes are not always trusted. Therefore, what the IAD needed to do as it formulated CSfC was take advantage of the good things that commercial products bring to the table while minimizing the disadvantages—and all with the end goal of composing commercial products together into a solution that is sufficient for protecting classified information (*CSfC overview* 2014).

CSfC Strategy

CSfC is founded on the principle that properly configured layered solutions can provide adequate protection of classified data in a variety of different applications. Unlike the GOTS arena where all of the security features are built into a single GOTS device, CSfC relies upon compositions of multiple commercial products to enable users to communicate securely based on open, commercial standards that distribute the risk over the entire architecture (*CSfC overview* 2014). Mrs. Plunkett said,

> "First and foremost, our whole strategy from a classified mobility perspective is very much focused on Commercial Solutions for Classified…It's our intent that we would deliver, end-to-end, a solution that is reliant on all commercial components and we believe we can do that". (American Forces Press Service 2012)

Strong Industry Partnerships

The IAD wants commercial component vendors to adopt open, non-proprietary standards for security and interoperability. In addition, the IAD is increasingly working in a collaborative manner with commercial industry to prototype end-to-end commercial solutions.

According to Mrs. Plunkett, '"We very strongly believe that, in the absence of our ability to be able to leverage the capacity of industry to deliver security and components that we need, we will not be able to meet the demand signals from our customers"'. Plunkett said NSA has had '"phenomenal partnership and cooperation across industry. Partnerships are critically important so [that] we establish that mobile enterprise capability"' (American Forces Press Service 2012).

Benefits of Adopting Open Standards and Commercial Products

There are no releasability issues with commercial equipment, and this makes CSfC attractive in satisfying coalition/diplomatic requirements where the delivered system could be outside the direct control of U.S. personnel (*CSfC overview* 2014).

Commercial devices are not considered Controlled Cryptographic Items (CCI); consequently, there are no special CCI handling caveats or restrictions associated with a CSfC solution. Customers apply existing information technology policies for safeguarding the processing of classified information (*CSfC overview* 2014).

There is also increased industry involvement in determining the security requirements for key technology areas (for example, Firewalls and Virtual Private Network (VPN) Internet Protocol Security (IPsec) Clients). The IAD does not levy security requirements unilaterally without the

benefit of input from and discussion with industry. Today, technology communities consisting of industry and academia work collaboratively together to identify the security requirements necessary for commercial components (*CSfC overview* 2014).

Perhaps the greatest benefit realized by customers is the speed with which solutions can now be delivered. This results in customers' being able to leverage commercial products while they are still current (*CSfC overview* 2014).

Built-In Assurance

Every CSfC solution has a number of designed-in assurance features. The trustworthiness of CSfC solutions and the components comprising them are validated by the IAD through a variety of methods:

- World-class cadre of IAD-trained systems security engineers (SSEs) design and develop CSfC security architectures;
- Senior NSA technical leaders review the architectures and provide oversight on all CSfC capabilities;
- Mitigations for vulnerabilities are applied based upon preliminary risk-assessment results according to National Institute of Standards and Technology (NIST) Special Publication 800-30;
- Customer registrations of CSfC solutions provide the IAD with situational awareness concerning which components are being used and where;
- Commercial components are tested against published U.S. Government Protection Profiles and internationally recognized Common Criteria security requirements;
- Commercial components that implement NSA-approved Suite B cryptographic algorithms are subjected to additional testing and validation of their implementation in accordance with NIST Federal Information Processing Standard (FIPS) 140-2; and
- Solution testing ensures compliance with NSA standards for the protection of classified information (*CSfC overview* 2014).

Extending Diversity Concepts

The IAD's approved CSfC Capability Packages include requirements calling for manufacturer diversity among specified component pairs contained within CSfC solutions. These requirements exist because the CSfC process seeks to ensure that the layers in a given CSfC solution are independent. Although different vendor logos do not imply independence, manufacturer diversity is often a good proxy that aligns with the intended requirement for independence among the layers. Understanding that dual-manufacturer requirements are imperfectly aligned with the independence goal, the IAD will be amending the current dual-manufacturer requirements (CSfC Program Management Office 2014).

In the near future, Original Equipment Manufacturers (OEMs) will have the opportunity to make a case that two of the OEM's product-lines are sufficiently diverse to be used as separate layers in a CSfC solution. Work is ongoing within the IAD to establish a process that will enable this significant change. Use of different codebases and diversity in the overall cryptographic implementation is among the characteristics that could contribute to the case for independence among distinct product lines made by the same OEM. Until a process is established for considering and adjudicating claims of independence among components made by the same

manufacturer, the current dual-manufacturer requirements remain in place (CSfC Program Management Office 2014).

Protection Profiles

USG Protection Profiles are specifications developed and published by the NSA Commercial Solutions Center (NCSC) National Information Assurance Partnership (NIAP) that capture security requirements for commercially produced technology. Protection Profiles are the mechanisms through which the USG conveys to commercial component vendors the security features needed to protect classified information. As of this writing, the NCSC has published 17 Protection Profiles, many of which directly support CSfC requirements. These published profiles cover technologies ranging from mobility wireless Local Area Networks (LANs) to applications and network devices. Fifteen other Protection Profiles are currently in development addressing a wide variety of technologies from system access controls to encrypted storage and virtualization (NSA/NCSC 2014).

Once a vendor builds a commercial component that meets the requirements of one or more Protection Profiles, that vendor submits the component to one of the authorized NIAP labs for compliance testing. This commercial lab testing is accomplished more quickly and inexpensively today than in the past due to significant transformations within NIAP to generate Protection Profiles written in precise, testable language (NSA 2014).

After successful NIAP testing, the commercial component is further qualified for use in a CSfC solution with the vendor signing a Memorandum of Agreement (MOA) with the NSA/IAD Business Affairs Office. The purpose of this MOA is to describe the vendor's and the IAD's responsibilities and obligations. After both parties sign the MOA, the commercial component is added to the CSfC Components List, which is published on NSA.gov. At this point, customers/integrators can select the component and use it in a CSfC solution. These approved commercial products on the CSfC Components List serve as the building blocks for CSfC (NSA 2014).

Capability Packages

Capability Packages serve as the foundation for NSA's commercial IA strategy. They provide general descriptions of the intended operational capabilities, along with the associated security architectures. Capability Packages contain vendor-neutral information enabling customers and integrators to implement their own solutions. Using the information in Capability Packages, customers/integrators make product selections from the CSfC Components List while following the specifications to create an architecture with specific commercial products configured in a particular manner (NSA 2014).

NSA prepares classified risk assessments that are available to appropriately cleared individuals. These risk assessments outline the specific residual risks related to implementing a Capability Package. These classified assessments explain the risks and demonstrate that secure commercial IA solutions are possible when customers adhere to the Capability Package's published specifications (NSA 2014).

Balancing Risk

Deciding what constitutes acceptable risk is always an important consideration for Approving Officials. Raising the bar in a security solution to mitigate vulnerabilities often comes at the expense of both development resources and operational efficiencies (Joyce 2012).Calculated risk-management decisions must occur with the Approving Official's eyes wide open. Commercial IA products must deliver a quality user experience that meets the requisite security policies without sacrificing operations (Joyce 2012).

The DIRNSA, as National Manager for National Security Telecommunications and Information Systems Security, under NSD-42, has authorized the IAD to develop and approve IA techniques to secure NSS. In executing this authority, the National Manager provides guidance regarding the appropriate combinations of commercial products. The published Capability Packages represent techniques that have been developed and approved by the National Manager as a commercial strategy suitable for protecting classified information. Certainly, this is predicated on the customer's implementation being configured and maintained as required by the Capability Package (CSfC Program Management Office 2014).

As with any new capability, customers must follow applicable certification and accreditation processes in order to obtain their Approving Official's authorization to run the solution on their operational networks. The customer's CSfC solution registration, which includes the Approving Official's assertion that the implementation is fully compliant with the Capability Package and that the residual risk is acceptable/mitigated, serves as the National Manager's acknowledgement of the customer's implementation of the CSfC solution (NSA 2014).

Today's CSfC Capabilities

There are several IAD-approved Capability Packages published on NSA's unclassified website at www.nsa.gov. These include Enterprise Mobility Capability Version 2.3, Campus Wireless LAN Version 1.0, and VPN Version 2.0 (NSA 2014).

Mobility

The NSA Mobility Program was established in response to the substantial and justified urgency for delivering Mobility solutions that securely provide the rich user experience of commercial technology. As customers and partners accelerate toward agile and mobile communications, the IAD has the responsibility for providing mobile capabilities that can evolve at the pace of today's commercial market, provide security, and enhance user experience (Mobility Program Management Office 2014).

The delivery of secure mobile capabilities to the USG and DoD requires a scalable approach to solution development. The first Mobility Capability Package was published in 2012 and encompassed the five major categories of the mobile ecosystem components: Secure Voice, Operating System/Applications and Mobile Device, Mobile Transport (Carrier), Mobile Enterprise Infrastructure, and Interoperability. The Mobility Capability Package will continue to evolve based on the IAD's analysis, prototyping, and piloting initiatives (Mobility Program Management Office 2014).

The current Mobility Capability Package is the initial release of the Enterprise Mobility Architecture for Secure Cellular. This version contains guidance on the required procedures necessary to build and implement a cellular voice and data capability using commercial-grade cellular mobile devices and infrastructure. Future releases will build on this architecture and will include mobile device management, international roaming, unified communications and enterprise services, and public key infrastructure. Ultimately, they will integrate the Wi-Fi service with an expanded list of end devices (NSA 2013).

Campus Wireless

The IAD published a Campus 802.11 Wireless LAN Capability Package to meet the demand for commercial tablet and laptop computers to access secure enterprise services over a campus wireless network. This provides the ability for customers to implement layered encryption between a classified network site and End User Devices. The Capability Package takes lessons learned from two proof-of-concept demonstrations, which included the layered use of commercial products for the protection of classified information (NSA 2014).

The solution architecture described within the Campus Wireless LAN Capability Package is supported by the use of wireless devices to access sensitive data and enterprise services while minimizing the risk when connecting to existing USG enterprise networks. Government-managed campus-area wireless networks provide controlled connectivity between mobile users and the broader USG enterprise. The term ‘campus’ refers to any area that is physically protected to the classification level of the classified network data. This physical area includes secure facilities and tactical environments when the physical controls are deemed appropriate by the Approving Official (NSA 2014).

Virtual Private Network (VPN)

To meet the demand for data in transit solutions, the IAD delivered a VPN Capability Package that enables customers to implement VPNs between two or more sites and VPNs between fixed sites and End User Devices. The Capability Package takes lessons learned from five proof-of-concept demonstrations, all of which included a layered use of commercial products for the protection of classified information (NSA 2014).

The VPN Capability Package describes a general VPN solution to protect classified information as it travels across either an untrusted network or a network of a different classification level. The solution supports interconnecting two or more networks operating at the same security level via a VPN. The VPN solution uses two nested, independent IPSec tunnels to protect the confidentiality and integrity of data as it transits the untrusted network. The two tunnels protecting a data flow are generated by VPN Gateways implemented as part of the network infrastructure or by VPN client software running on an End User Device (NSA 2014).

The solution also supports connecting individual End User Devices to a network via the VPN solution if the device and the network operate at the same security level. This solution architecture provides sufficient flexibility to be applicable to many use cases of VPN implementations (NSA 2014).

Early CSfC Adopter

One of CSfC's early adopters is a Combatant Command that has welcomed CSfC at its site. The command's mission mandates it exchange information and coordinate with foreign partners. According to the CIO, the command was looking for an easier way of achieving the capability previously provided by proprietary GOTS encryption solutions. After weighing the benefits offered by CSfC, the command decided to use the VPN Capability Package (*CSfC* 2013).

The CSfC VPN solutions have benefitted this command in a variety of ways. The command has been able to cut back on the amount of travel required to sustain the solutions. That results in saving both cost and time. According to the command, these systems are proving to be easier to sustain than their GOTS predecessors, resulting in additional time savings. Overall, the CSfC implementations throughout the command have provided a tangible return on investment, and the command will look to CSfC in the future for additional capabilities (*CSfC* 2013).

Exciting Times for CSfC

Five thousand users across nearly four dozen sites including civil and military agencies are using, preparing to use, or prototyping composed commercial IA solutions. For the first time, customers are registering CSfC solutions to protect classified information, and this signals a significant transformation for the IAD as customers begin migrating to commercial IA solutions (CSfC Program Management Office 2014).

CSfC has a presence on NSA's public website (http://www.nsa.gov/ia/programs/csfc_program). Capability Packages are published providing vendor-neutral specifications to many different customers who are all trying to solve the same problems. NSA has held several 'Industry Days' and military service 'IA Days' to socialize CSfC with client stakeholders and commercial supplier communities. In the last year, the IAD has interacted with more than 1,000 customers, policymakers, integrators, and suppliers to continue promoting acceptance of CSfC. Feedback from many of them has confirmed that CSfC is welcomed as the right program at the right time (CSfC Program Management Office 2014).

NSA has established a Mobility Innovation Center; delivered an NSA laptop pilot; and published specifications for mobility, wireless, and VPN architectures. For secure cell phones, NSA operated the FISHBOWL pilot which was approved for TOP SECRET voice and data. The IAD is currently working on an operational capability based on this pilot with DISA as a partner. Additionally, the IAD has run pilots using consumer tablets as both e-readers and connected devices (Mobility Program Management Office 2014).

Despite these current strides and successes, Debora Plunkett's focus is forward: '"Looking ahead ... we've got to make sure we're constantly looking at the user experience [and] responding to the needs of the user. We continue ... to prototype and pilot different services"' (American Forces Press Service 2012).

Trusted System Integrators

Trusted system integrators are instrumental to CSfC's success because they will serve as the honest brokers for decision makers who will be responsible for approving the use of commercial IA solutions at their sites. To ensure there is a cadre of qualified trusted integrators, IAD hosted a

CSfC Integrator forum in November 2013. Representatives from sixteen commercial companies/military service labs attended to discuss the issues faced by system integrators when composing commercial IA solutions (CSfC Program Management Office 2014).

The IAD will consider organizations as trusted CSfC integrators if they are able to demonstrate the following capabilities:

- Meet the IAD published criteria,
- Assemble/integrate commercial components according to CSfC Capability Package requirements,
- Perform system security testing on the commercial IA solution to verify all Capability Package security requirements are addressed,
- Develop the required technical body of evidence for submission to the relevant Approving Officials, and
- Deliver ongoing lifecycle support for the solutions the integrator develops if required by the customer (CSfC Program Management Office 2014).

Improved Dialogue

Through a variety of communication channels, the IAD is promoting CSfC's core message that properly configured, layered commercial solutions are sufficient for protecting classified information. The objective of CSfC's communication strategy is to promote consistent, timely, and accurate information to targeted audiences that include congressional stakeholders, policymakers, government customers, and commercial industry (CSfC Program Management Office 2014).

The IAD has established multiple web presences across UNCLASSIFIED, SECRET, and TOP SECRET networks in order to maximize the dissemination of key CSfC artifacts (for instance, Capability Packages). It has hosted numerous 'NSA Days' and 'IA Days' with senior leadership across Combatant Commands, military services, and federal agencies to provide opportunities to promote awareness of CSfC capabilities. 'Industry/Integrator Days' with commercial product developers have proven to be effective ways to share the NSA's commercial-strategy vision and to obtain feedback from the private sector. The IAD is also a frequent exhibitor at Cyber/Security expositions and trade shows as these provide for direct interaction with commercial vendors and government customers, which contributes to the IAD's ability to socialize and reinforce key CSfC principles (CSfC Program Management Office 2014).

For the first time, DoD/Intelligence Community organizations are implementing NSA-approved commercial security solutions to protect their classified information. The IAD-approved Mobility, VPN, and Campus Wireless LAN Capability Packages describing how to implement commercial security solutions are now available to government customers, commercial product vendors, and commercial solution integrators (CSfC Program Management Office 2014).

Now Is the Time for Commercial Vendors to Get Involved

The NSA expects commercial industries to get their components certified and approved by NIAP and to have a signed MOA with the IAD. This is the path for getting commercial products eligible for use in commercial IA solutions. Industry now has access to the NSA-approved component-level specifications via USG Protection Profiles and system-level specifications via

Capability Packages. These specifications provide industries with the information they need in order to build commercial IA components and integrate the resulting layered commercial IA solutions (CSfC Program Management Office 2014).

The IAD expects industry to adopt open (non-proprietary) standards for interoperability and security. The following list provides a sample of some of these commercial standards.

- IPSec – The protocol suite for secure Internet Protocol communications. It authenticates and encrypts each IP packet of a communications session. IPSec also includes protocols for establishing mutual authentication and negotiating cryptographic keys for a session.
- Suite B Cryptography – Far more than just Advanced Encryption Standard (AES) for encryption, Suite B also includes cryptographic algorithms for key exchange, digital signatures, and hashing.
- Session Descriptions (SDES)-Secure Real-Time Transport Protocol (SRTP) – SDES Protocol Security Description for Media Streams is used to negotiate the encryption key for the Secure Voice over IP (VoIP) application. SDES-SRTP is used as a base protocol for true end-to-end key exchange/security.
- *NIST 800-164 guidelines on hardware-rooted security in mobile devices* (Draft) is a special publication that defines security fundamentals and capabilities and identifies a baseline for secure technologies (CSfC Program Management Office 2014).

How to Get Involved

NSA-approved commercial IA solutions are in use today to protect classified data. Because of NSA's commercial strategy, new opportunities have been created for increased competition in the commercial marketplace. By leveraging current technology and open standards, NSA makes it possible for everyone who uses commercial IA products to communicate securely using a wide variety of commercial devices (CSfC Program Management Office 2014).

Capability Packages may be downloaded by visiting the CSfC website at www.nsa.gov. To download the USG/Intelligence Community, customers can inquire at 410.854.4790. DoD, Military, and Combatant Command customers can inquire at 410.854.4200. Industry inquiries can be directed to the IAD Business Affairs Office at 410.854.6091 or bao@nsa.gov. General inquiries can be directed to the NSA IA Service Center at niasc@nsa.gov (*CSfC* 2014).

Conclusion

Although NSA's strategy for protecting classified information continues to employ both commercially based and traditional GOTS IA solutions, the IAD is looking first to commercial technology and solutions in order to provide its diverse customers with the flexibility they require to design, build, and implement commercial IA solutions that satisfy their security policies and requirements (*CSfC* 2014).

The IAD's customers support CSfC because it enables them to get IA solutions deployed faster without many of the challenges that typically come with proprietary GOTS solutions. Policymakers are encouraged by CSfC given their longstanding confidence in the IAD's reputation of protecting information and for outmaneuvering adversaries. Industry praises the IAD's 'Commercial First' model because it fully leverages the energy and innovation of the private sector (*CSfC* 2013).

References

American Forces Press Service 2012, 'NSA looks to industry for secure mobile capabilities', July.

Confidence in cyberspace 2013, video, NSA/IAD, November, FGGM.

CSfC 2013, video, NSA, TAKE2 video, August, CSfC PMO.

CSfC 2014, brochure, NSA/IAD, January.

CSfC overview 2014, briefing, NSA/IAD February.

CSfC Program Management Office 2014, NSA/IAD, February.

Joyce, R 2012, 'Bad user experience: another form of vulnerability' *Federal Technology Magazine*, April, viewed 9 April 2012, <http://www.fedtechmagazine.com/article/2012/04/bad-user-experience-another-form-vulnerability>.

Mobility Program Management Office 2014, NSA/IAD, February, <http://www.nsa.gov/ia/programs/mobility_program/index.shtml.>.

NSA 2013, *Mobility*, viewed 13 August 2014, <www.nsa.gov/ia/programs/mobility_program>.

—2014, *CSfC*, viewed 20 February 2014, <www.nsa.gov/ia/programs/csfc_program>.

NSA/NCSC 2014, *NIAP*, March, viewed 13 January 2014, <www.niap-ccevs.org>.

Using Classified Intelligence to Defend Unclassified Networks

N Ziring, B Thomas

Information Assurance Directorate
National Security Agency, Fort Meade, Maryland, United States
E-Mail: JIWfeedback@nsa.gov

Abstract: *Intelligence services, such as the National Security Agency, have access to unique information about adversarial cyber-exploitation and -attack capabilities. Nations such as the United States should be employing this unique but sensitive information in the defense of national security, government, critical infrastructure, and other networks, but doing so may expose the sources and methods behind the intelligence. Once exposed, access to that unique information may be lost. This paper describes the dilemma, presents a partial taxonomy of use cases for which solutions are needed, and offers avenues for supplying those solutions. In particular, solutions to the problem of using classified intelligence for defense of unclassified networks fall into three approaches. Properties and examples for each approach are presented, and advantages and disadvantages discussed.*

Keywords*:* *Network Defense, Cyber Security, Applications of Intelligence, Trusted Computing, Private Information Retrieval*

Introduction

Cyber-threat actors and cyber defenders engage in a continual competition, both improving their technologies, tools, and tradecraft in response to the other's successes. In the competition, the attacker has great freedom of action and can enjoy significant advantages when the defender lacks insight into the adversary's plans and intentions (Hutchins, Cloppert & Amin 2011; Williams, Shimeall & Dunlevy 2002; King, Orlando & Kohler 2012).

Intelligence services use specialized**,** and often clandestine**,** sources and methods to gather information about their targets. In the cyber realm, intelligence targets may include a variety of cyber-threat actors (usually foreign nations or trans-national groups)**,** and the intelligence gained may include critical details about a threat actor's software, command-and-control, and other assets. The United States (U.S.) intelligence and national security communities face a dilemma: intelligence-derived information, if used properly, can help shift the advantage back to the defender, but it is also highly sensitive and correspondingly classified. To employ the information in network defense exposes it to leakage or theft, particularly by cyber-threat actors. Once such information is compromised, it loses much of its unique defensive value because actors can modify their plans or adjust their tradecraft. More importantly, compromise of defensive information based on intelligence can also compromise the sources and methods used to obtain it, thus eliminating future intelligence.

Therefore, the U.S. national security community needs mechanisms that will allow effective use of sensitive intelligence information for defense of a broad range of networks, but which can also protect that information from unauthorized access, leakage, and theft.

Scope

The discussion and solution approaches in this paper are focused on intelligence information useful for directly detecting actions against a defended network or system. This includes the kinds of artifacts that network defenders derive today from forensic analysis of successful attacks, such as file and network signatures, communication addresses, and credentials. Higher-level intelligence, such as threat-actor goals and intentions, can also be very useful for defensive strategy, but are outside the scope of this paper.

General properties of solutions

Solutions for using classified intelligence information in defense should exhibit three main properties.

1. Defensive Effectiveness: the solution should meet the same criteria for accuracy, timeliness, and flexibility that apply to conventional defensive techniques. For intrusion detection, such criteria are well understood (Amoroso 1999; Scarfone & Mell 2009).
2. Confidentiality: the solution must protect the content of the intelligence information, and ideally even its existence, from all unauthorized parties.
3. Operational Relevance: the solution must be usable in one or more operational use cases (see Use Cases section below).

Defense is more than detection. An authorized operator who receives detection information based on classified information might choose to take a response action or configure a system to take action. In many cases, taking such an action will allow an external observer, or the attacker, to learn something about the intelligence information. For example, if the intelligence included a malicious document that an attacker planned to send, then blocking it could reveal to the attacker that the defender had advance warning about that document. The authorized defender must choose a response action based on risk management and cost/benefit analysis. But a critical aspect of the confidentiality property is that the solution must protect all that it can—prior to defender action, the attacker and other observers must be able to gain no knowledge of the intelligence information. Even after a defensive response, the exact content must remain secret. Confidentiality about defensive information can provide a powerful advantage for the defender. Means for maintaining that advantage, even in very hostile conditions, are the focus of this paper.

Related Work

The dilemma identified in the Introduction presents a specific case of a more general problem: using information for a task without revealing that information. This section provides an overview of some of the considerable research and development devoted to solutions for specific problems in this vein.

Private Information Retrieval (PIR) and privacy-preserving queries are very active research areas. The usual problem involves extracting information from a database or stream of data, without revealing the query or extraction criteria. Many of the proposed solutions apply cryptographic techniques. Kushilevitz and Ostrovsky demonstrated computationally private database retrieval in

1997. Other early work on PIR was performed by Chor, Goldreich, Kushilevitz, and Sudan (1998). Boneh, Gentry, Halevi, Wang, and Wu have applied partially homomorphic encryption (encrypted computation) to private database queries (2013), and homomorphic techniques have also been applied for secure pattern matching (Yasuda *et al.* 2013). Surveys of PIR work are available in a general survey in 2004 (Gasarch) and a more focused database PIR survey in 2007 (Otrovsky & Skeith).

Privacy-preserving cryptographic techniques have been applied directly to network defense research. A classic problem in that area is correlation of security events while preserving the privacy of individual events. Lincoln, Porras, and Shmatikov reported on a privacy-preserving system for sharing alerts (2004). Ma, Chen, and Li presented an approach for privacy-preserving alert correlation (2010), and Li, Liang, Lu, Shen, Lin, and Zhu proposed applying privacy-preserving techniques to aggregation of critical infrastructure monitoring data (2012). Most recently, Niksefat, Sadeghiyan, Mohassel, and Sadeghian reported on a privacy-preserving intrusion detection system built on secure 2-party computation (2013).

There has also been substantial research and development on computational platforms that can isolate or protect certain computations, ensuring integrity and/or confidentiality. Early work on this topic was performed in the 1970s, for example, on virtual machines for isolation (Madnick & Donovan 1973). NSA researchers reported on NetTop™, which used a commercial hypervisor to support execution of multiple classification levels of processing, mutually isolated, on a single host (Meushaw & Simard 2000). More recently, frameworks for isolating and protecting specific computations and processes within a host without using virtualization have been codified by industry consortia (Global Platform 2011) and by particular companies (ARM Ltd 2009; McKeen *et al*. 2013).

Sometimes, the cryptographic and isolated computation approaches have been combined. For example, Wang and colleagues reported on implementing PIR on trusted hardware (Wang *et al*.).

Use Cases

There are many situations in which sensitive or classified information may be applied in cyber defense. This section presents a partial taxonomy of such situations, which may be used to assess the suitability of a particular solution to candidate applications. To illustrate the taxonomy, several specific use cases**,** which have arisen in practice**,** are presented and categorized.

Application taxonomy

The following factors are important to describing and understanding the network defense use cases for the application of classified information in the defense of unclassified systems.

The use cases differ greatly depending on the existence of ‘connectivity’ between the classified systems where the classified information originates, and the unclassified systems to be protected. Even when connected, the ‘bandwidth’ between these systems plays a major factor in the solution space. ‘High-bandwidth’ connections (that is, well connected) offer the most flexibility. ‘Low-bandwidth’ connections, such as those available to Forward Operating Bases (FOBs) in combat zones, are more challenging. The most challenging case is the defense of ‘disconnected’ unclassified systems that have no connectivity back to any classified systems.

To deal with the challenges of low-bandwidth connectivity, security assessments or incident responses can be conducted locally by deployed, cleared security personnel or remotely, albeit slowly. In the case of disconnected systems, there is little choice but to deploy cleared security personnel to conduct a local assessment. In both cases, the cleared security personnel will have access to classified information to use during the security assessment or incident response. The various use cases are summarized in **Figure 1**.

	Connectivity		
	High-Bandwidth	Low-Bandwidth	Disconnected
Local operations		Use Case 2	Use Case 4
Remote operations	Use Case 1	Use Case 3	

Figure 1: Use Case Summary

Use cases

The following four use cases are described using the factors outlined above in the taxonomy, sorted roughly from easiest to hardest.

Use case 1: high-bandwidth remote monitoring

In this use**,** connectivity exists from the classified systems to the unclassified systems to be protected via appropriate cryptographic protection mechanisms and cross-domain solutions (for instance, guards). The security operators work on classified networks and have full access to classified adversarial indicators. They do not need to travel or deploy to the field locations where the unclassified systems reside**,** and they are able to manage remotely the sensor suites that implement the detection capabilities. The core challenge is to enable the sensor suite to search for adversary indicators on classified or unclassified networks, using sensitive intelligence information, without exposing that information to compromise. Typically the on-site IT staff and the Computer Network Defense (CND) Service Providers (CNDSP) do not have access to the remotely managed sensor suites. **Figure 2** illustrates the structure of Use Case 1.

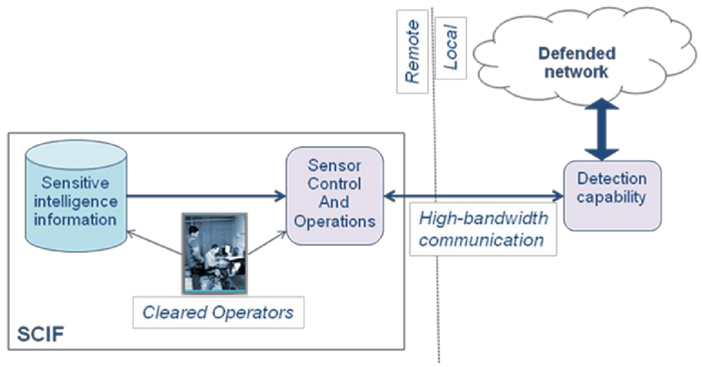

Figure 2: Use Case 1: High-bandwidth remote monitoring

Use case 2: low-bandwidth local assessment
In this use case, connectivity exists from the classified systems to the unclassified systems to be protected**,** but that connectivity is low-bandwidth. Defensive-operations personnel are deployed to the field location**,** and may have access to secure communications**,** but are limited to low-bandwidth. For example**,** they often have modest classified communications access, such as e-mail or web, via Joint Worldwide Intelligence Communications System (JWICS) and/or Secret Internet Protocol Router Network (SIPRNet). There is usually no direct communication between classification levels other than the team members themselves or media, which can be moved manually (for instance, CD-Rs). **Figure 3** illustrates the low-bandwidth local assessment use case.

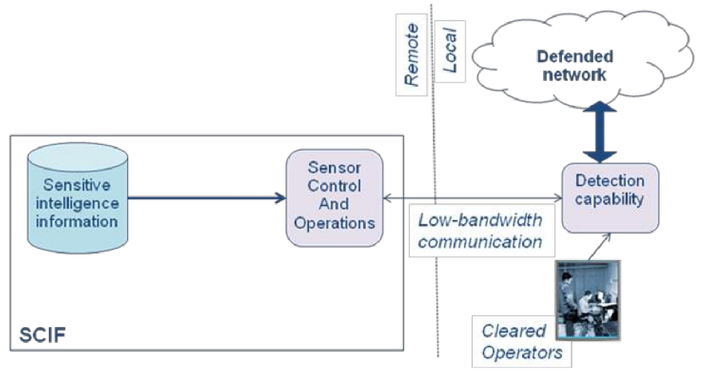

Figure 3: Use Case 2: Low-bandwidth local assessment

For example, NSA defensive operators are cleared and have the training and authorization to handle sensitive intelligence data. But they conduct a wide variety of operational assessments and incident responses. Every evaluation is different: the degree of access, convenience, timeliness, and bandwidth all vary. Some sites may have only unclassified or only secret and unclassified networks. When the operation involves deploying a team to a remote site, the team members typically cannot carry sensitive intelligence information to the site, because most sites lack Sensitive Compartmented Information Facility (SCIF) facilities. The goal for the assessment team is to use its physical presence**,** and deep access to site assets**,** to search for adversary traces on networks, servers, and workstations at all classification levels (especially unclassified).

Use case 3: low-bandwidth remote monitoring
In this use**,** connectivity exists, at low-bandwidth, from the classified systems to the unclassified systems to be protected. The security analysts work on classified networks and have full access to classified adversarial indicators. They do not need to reside at or deploy to the field locations where the unclassified systems reside. They are able to manage the sensor suites remotely. There are two core challenges. The first is shared with the use case 1, which is the challenge to enable the sensor suite to search for adversary indicators on classified or unclassified networks, using classified adversary information, without exposing that information to possible compromise. Secondly performance is greatly reduced due to the low-bandwidth**,** which makes even simple tasks, such as scanning a network and returning the results, difficult and slow. Typically the on-

site IT staff and the CNDSP staff have little, if any, access to the advanced detection capabilities. **Figure 4** illustrates the structure of this use case.

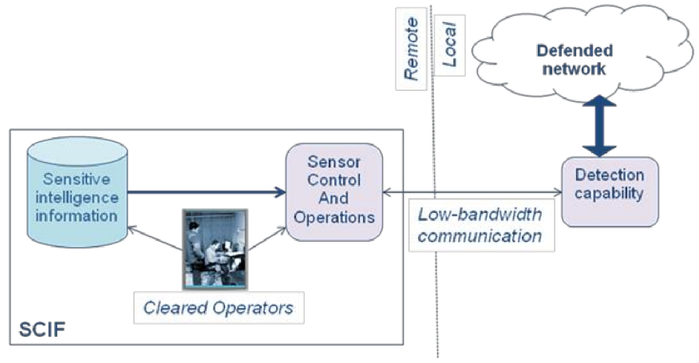

Figure 4: Use Case 3: Low-bandwidth remote monitoring

Use case 4: disconnected incident response

In this use case, connectivity does not exist between the classified systems and the unclassified systems to be protected. Security personnel must be deployed to the field location to provide incident response services. There is no direct communication between classification levels other than the knowledge and experience of the team members. The team members have no secure communications, and they cannot carry classified information to such sites, due to lack of a SCIF. The goal for the team members is to use their physical presence and deep access to site assets to search for adversary traces on networks, servers, and workstations at all classification levels (especially unclassified). **Figure 5** illustrates the disconnected incident response use case.

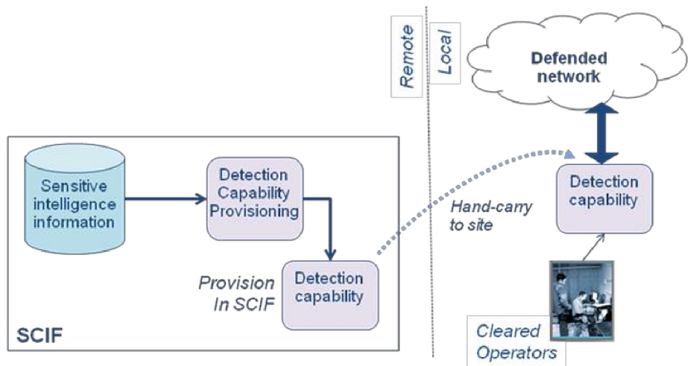

Figure 5: Use Case 4: Disconnected incident response

Solution Approaches

This section describes three approaches for addressing the dilemma of using sensitive information in cyber defense, while meeting the general requirements discussed in the introductory section of the paper. The approaches may be characterized roughly on the primary mechanisms that protect the sensitive information: physical isolation, software (computation) isolation, and cryptographic

isolation. Each approach has advantages and short-comings. An important consideration for solutions is whether the approach works well for network-intrusion detection, host-intrusion detection, or both.

Physical-isolation approach: secure appliances

A traditional means of protecting a secret is to lock it in a strong box. This means can be used for protecting intelligence information while using it for defense, by building a defensive appliance as a discrete device (for example, a network sensor). In this model, the sensitive information stays in the box and is processed against network data. Such an appliance can operate in-line at the boundary between two networks, or via a passive tap that copies information from one or more network domains to the appliance (Scarfone & Mell 2009). Secure appliances must have the following properties:

- They must support access control and secure communications for secure use by authorized parties.
- They must accept input from the defended systems or networks, and use the intelligence information to detect malicious events or behaviors.
- They must notify authorized parties of detection in a timely and secure manner. They do not reveal the intelligence information or the detection notifications to any unauthorized parties.
- They are highly resistant to physical tampering and other forms of close-access attack. In particular, the sensitive information should be erased if attempts at physical attack are detected (NIST 2001).

Figure 6 shows how a secure appliance acts as an extension of the classified systems, operating attached to the local unclassified network. It is controlled and managed solely by cleared operators.

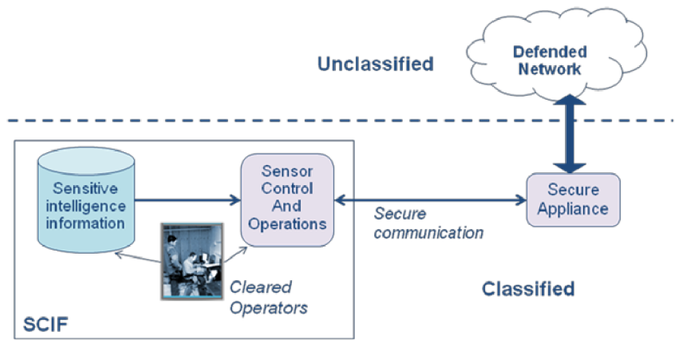

Figure 6: Employing a Secure Appliance

Most commercial network appliances, such as intrusion-detection systems, satisfy properties 1-4 to some degree. To be considered as sufficiently secure to protect classified intelligence, the mechanisms used must meet relevant standards (see CNSS 2012). Only a few commercial network devices are built with these properties. To illustrate this approach, two devices designed with all five properties in mind are described as examples below.

Example 1: Cloud Shield CS-4000 trusted network security platform
The CS-4000 is a commercial Deep Packet Inspection (DPI) network device. It incorporates features for protecting the signatures and detection methods loaded into it, protecting communication with authorized users, and resisting physical attacks (Cloud Shield 2013). It implements IPSec with Suite B-compliant cryptography for communication, giving cryptographic strength of function sufficient to protect classified information (CNSS 2012).

Example 2: KG-175G trusted sensor
The KG-175G is a commercial network sensor and network encryptor, built specifically to satisfy U.S. Government requirements for protection of classified information. It uses network packet signatures to detect particular activity on a network link, and protects the signatures inside its tamper-resistant chassis. The KG-175G implements the U.S. government standard HAIPE network-layer encryption protocol for protecting communication between the appliance and authorized users; its implementation of HAIPE is certified by NSA as sufficient for protecting classified information (General Dynamics 2013).

Assessment of secure-appliance approach
Both examples share some features that help them support the use cases outlined in section 3. First, sensitive information is stored inside the appliance during operation, within the tamper-resistant boundary of the device. Second, the device must have long-term cryptographic credentials, which are stored internally, cannot be exported, and are erased if physical tampering is detected. Third, the device uses the credentials to establish cryptographically secure communication with cleared, authorized operators (who must possess a credential that the device has been configured to trust). Fourth, notifications of activity detection and management functions can only occur over the secure link.

The secure appliance approach can satisfy all three general properties for solutions to a substantial degree. Defensive effectiveness will depend on the features of the particular appliance; but as long as the defensive signatures and detection methods are downloaded into the appliance, then it should be able to perform the same operations as other network defense appliances. Confidentiality will depend on three factors: 1) assurance provided by the secure communications between the appliance and the cleared operators; 2) resistance to close-access attacks, including physical anti-tamper and integrity; and 3) software integrity and resistance to attacks from the defended network. Operation relevance can be judged by the breadth of operational use cases that the secure appliance can support. This combination of features allows the example secure appliances to support use case 1, well-connected monitoring, very effectively, and should also be usable for use cases 2 and 3. But secure appliances such as the examples above are not suitable for use case 4, disconnected incident response, because they require real-time connectivity to a SCIF for reporting results.

Because the secure appliance is a separate device on the network**,** it is most suited to performing detection against network activity; applying this approach to detecting malicious activities on individual hosts would require sending all relevant activity from the hosts to the appliance. While this is possible, it does not match current practice for host-intrusion detection and might present scaling issues for large collections of hosts.

Computation-isolation approach: secure execution environments

Another approach to supporting the use of sensitive intelligence is to perform detection computations within a trusted, isolated compute environment or 'secure enclave'. The enclave must provide assurance of three critical properties: 1) sensitive information cannot leak out to, or be extracted by, unauthorized parties; 2) only authorized parties can configure or add new sensitive information to the enclave; and 3) the software used to perform the detection is the authorized, correct software, which can be confirmed by external authorized parties.

In general, a secure enclave acts as a trusted proxy that can act on behalf of cleared defenders, performing detection using the sensitive information, but isolating the information, computations, and results from any other processing or users on the platform. To support defensive-use cases, the secure environment must also have some means to support secure communication with the cleared defenders, to accept new detection settings**,** and to deliver results. **Figure 7** shows how the secure enclave works.

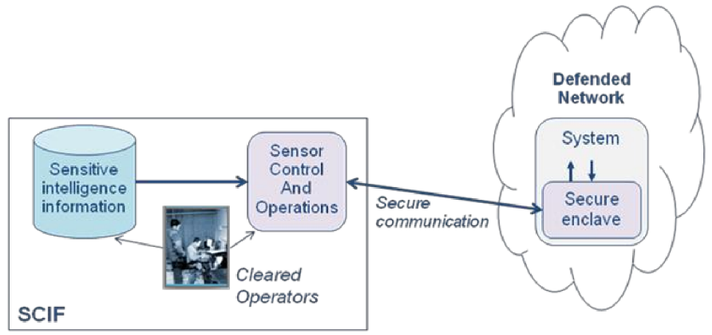

Figure 7: Employing Computation Isolation

Many different mechanisms have been proposed and studied for creating trusted, isolated computation environments; the topic has a long history, for example, see Madnick and Donovan (1973) and later Meushaw and Simard (2000). Many modern processors are built with hardware virtualization or domain separation features that can isolate computation and data (for instance, ARM TrustZone (ARM Ltd, 2009)), which can then be leveraged to satisfy properties 1 and 2. Until recently, there was no integrated hardware system that directly supported all three properties, but Intel Corporation recently announced processor features that do.

Example: Intel Software Guard Extensions™

Software Guard Extensions (SGX) allow an application to create a protected container—an enclave—which is a protected area in the application's address space that provides integrity and confidentiality assurance (McKeen *et al*, 2013). SGX integrity and confidentiality guarantees are enforced by the processor, even against fully privileged software running on the same processor. Further, an SGX enclave has cryptographic keys that allow it to store data encrypted only for itself. Together, the execution and storage features allow an SGX enclave to satisfy property 1. Launching an enclave initiates a process of measurement which allows the enclave to attest to its integrity to a remote party and securely exchange information with those parties (Anati *et al*. 2013). These features allow an SGX enclave to satisfy properties 2 and 3.

Assessment of computation isolation approach

A secure-compute-environment implementation that satisfies the three properties above can support all four use cases. The sensitive data is accessible only to the software running in the enclave, but the enclave can accept data from its local host and/or network environment. Detection computations using the sensitive data proceed in isolation, and results stored are protected. For use cases 1 and 3, results are reported to the remote cleared operators. For use cases 2 and 4, results can be reported locally only to authorized users, or wrapped for later unwrapping in the SCIF. For use case 4, in particular, sensitive information can be provisioned into the enclave before the incident response.

When used for network detection, the enclave simply acts like the secure appliance. For host detection, the enclave runs on the host, protecting the sensitive information but performing computation against local host conditions and behavior. Note that, in both cases, the local software environment can prevent the secure enclave from detecting behavior simply by corrupting the information passed to it, but it cannot extract the sensitive information or masquerade as the trusted computation to external authorized parties.

Cryptographic-isolation approach: encrypted computation

The cryptographic-isolation approach protects sensitive information by provisioning it to the unclassified environment only in encrypted form, performing cryptographically masked computations, and returning an encrypted result to the classified environment. Inside the SCIF, the sensor-control system can decrypt the results and inform the cleared operators of any positive detection events. Unauthorized parties in the unclassified environment see only cipher text; if the cryptography is strong, then they can gain no knowledge about the sensitive information. **Figure 8** (below) shows a simplified overview of how the detection computations would work. This form of encrypted computation for intrusion detection is Private Information Retrieval (PIR) (Chor *et al*. 1998); in particular, it is single-database PIR (1dPIR), where the state of the unclassified environment, or network traffic passing through it, is the 'database' from which a result must be privately retrieved.

For example, a PIR scheme suitable for protecting sensitive information and operating as a sensor requires that a system in the unclassified environment be provisioned with a public key; the associated private key is held back in the SCIF. To deploy a particular set of sensitive information to the unclassified environment, the cleared operators use the private key to encrypt it, and send the encrypted block (the 'query') to the detector. For each set of local state data (for instance, events, packets, file properties) in the unclassified environment, the detector computes an encrypted result from the dataset using the encrypted query and the public key, and returns the encrypted result to the SCIF, as shown in **Figure 8**.

There have been a number of cryptographic algorithms described in the literature that can support PIR in the way shown in **Figure 8**. They offer a variety of detection power and efficiency trade-offs in terms of computational overhead and the size of the encrypted query and result. (Note: most such systems are 'partially homomorphic', meaning that the encrypted computation can only include certain operations. A 'fully homomorphic' system can perform any computation; only a few such systems have been reported, and all impose high computational overhead which limits their capacity.)

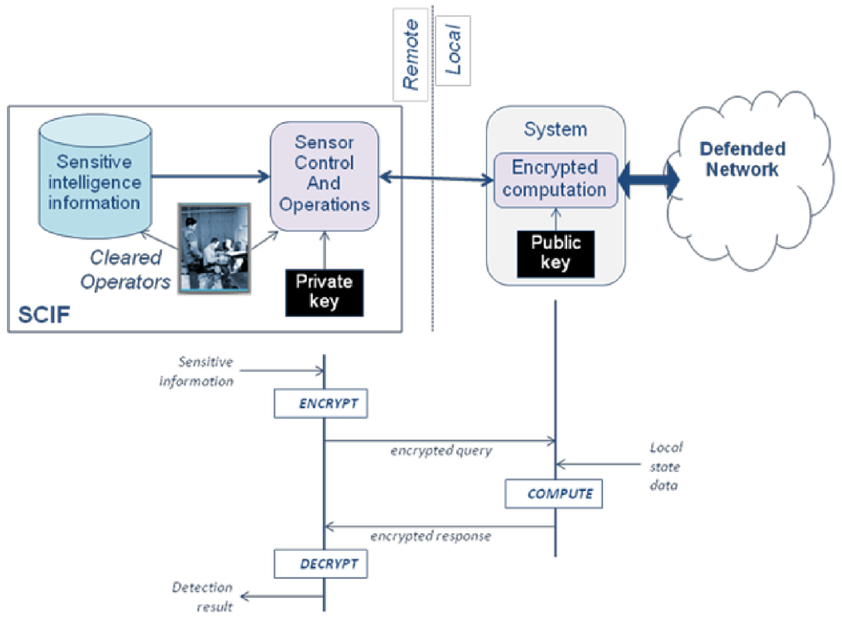

Figure 8: Employing Cryptographic Isolation

Example: using single-database PIR

The single-database PIR scheme described in Chang (2004) allows the sensor to store a database of N strings, and the detection system in the SCIF to retrieve for any single string of index $n < N$, without revealing n to the sensor. It uses the Paillier encryption system and supports efficient communication between the deployed sensor and the control system in the SCIF. By using a hash bucket scheme, this fundamental 1dPIR functionality can be turned into a private hash table lookup. For example, if the sensitive information being queried were the MD5 hash value of a file, the sensor system could group the MD5 values into a set of N buckets using, perhaps, the 16 bit prefix of each value. Using the PIR scheme, the control system in the SCIF could retrieve the subset of MD5 values matching a given prefix and search that subset for the sensitive value, without revealing even the prefix to the sensor.

Assessment of cryptographic-isolation approach

Because the private key must not leave the SCIF, the cryptographic-isolation approach is most suitable for use cases 1 and 3. Depending on the efficiency of the particular PIR scheme selected, application to use case 3 must be assessed to see whether available bandwidth will support necessary detection performance. For use case 4, a system could be provisioned with encrypted queries, and the encrypted responses could be decrypted when the system is returned to the SCIF. The cryptographic-isolation approach works for both host and network detection.

At present, the homomorphic encryption techniques used for encrypted computation impose substantial computational overhead. Even recent work specifically optimized for intrusion detection shows significant latency (Niksefat *et al*. 2013). Further research and optimization will

be required before cryptographic-isolation approaches will offer sufficient performance for general-purpose intrusion detection, but they may be adequate for specialized niche applications.

Directions for Future Work

All three approaches offer practical means to perform detection for sensitive classified data in an unclassified environment, while protecting the sensitive data. But only the secure-appliance approach is currently realized in available implementations.

For the computational-isolation approach, a small number of potentially viable technology platforms are available. Initial research must focus on building a variety of host- and network-detection facilities on top of these platforms, and characterizing their security, functionality, and performance. The results can be used to improve the platforms, and to build commercial implementations for government defense pilot programs.

For the cryptographic-isolation approach, numerous PIR schemes have been published that can support various detection scenarios. More research is needed to characterize the computation and communication overhead of these schemes in realistic network-defense scenarios. Next, developers need to integrate the PIR schemes into sensors and their control systems, to make the cryptographic assurance they offer available for government and private sector operations.

Conclusions

Classified information can provide a unique benefit to defense of unclassified networks, but only if the confidentiality of that information (and the source and methods behind it) can be protected. Three basic approaches are being explored in the research and development community for supporting this. **Figure 9** (below) summarizes those three basic approaches and their current maturity.

Approach	Security Foundations	Applicability	Maturity
Secure Appliance	Physical isolation, secure communication path	Use cases 1-3, Network, broad functionality	High – products available
Secure Execution Environment	Software isolation, secure communication path	Use cases 1-4, Network and host, broad functionality	Medium – technology available, products not yet available
Encrypted Computation	Cryptographic isolation	Use case 1-4, Network and host, limited functionality	Low – algorithms available but may not scale; toolkits and products not yet available

Figure 9: Summary of Solution Approaches

At the moment, only the secure-appliance approach is ready for operational use, and the range of products specifically designed for protecting classified defensive information is modest. The computation-isolation approach shows great promise. It appears to be the most flexible of the three approaches, and relevant secure execution features are being offered on modern processors. However, no products or applications that employ this approach are yet available for defenders to deploy. The encrypted-computation approach has the potential to offer very high assurance of

confidentiality because the classified information is never exposed unencrypted on the defended network. But limitations imposed by performance and functionality constraints make it the least flexible of the three.

By deploying secure appliances, it is possible today to use classified intelligence information for defense of unclassified networks, as long as suitable communications with cleared operators can be provided. It is not yet feasible to employ classified information for defensive operations on unclassified hosts. The computation-isolation and cryptographic-isolation approaches offer means to support hosts, but further development of products and applications is needed. In the case of encrypted computation, further research is ongoing to improve performance and breadth of functionality.

As these technologies continue to develop and become available, the U.S. intelligence and cyber-defense communities must create the policies and practices to take advantage of them. This will give the U.S. and its allies a powerful ability to apply our significant intelligence capabilities to defending our networks**,** while managing the risks such application could pose to those capabilities.

References

Amoroso, EG 1999, *Intrusion detection: an introduction to Internet surveillance, correlation, trace back, traps, and response*, Intrusion.Net Books, n.p.

Anati, I, Gueron, S, Johnson, SP &Scarlata, VR 2013, *Innovative technology for CPU based attestation and sealing*, white paper, Intel Corporation, Beaverton, Oregon, United States.

ARM Limited 2009, *ARM security technology: building a secure system using TrustZone technology*, white paper, PRD29-GENC-009492C, April.

Boneh, D, Gentry, C, Halevi, S, Wang, F & Wu, DJ 2013, 'Private database queries using somewhat homomorphic encryption', *Proceedings of the 11th international conference on Applied Cryptography and Network Security*, ACNS, Springer-Verlag, Berlin, Heidelberg, Germany, pp. 102-18.

Chang, YC 2004, 'Single database private information retrieval with logarithmic communication', *Information Security and Privacy*, Springer, Berlin, Heidelberg, Germany, pp. 50-61.

Chor, B, Goldreich, O, Kushilevitz, E & Sudan, M 1998, 'Private information retrieval', *Journal of the ACM (JACM)*, vol. 45, no. 6, pp. 965-81.

Cloud Shield 2013, *CS-4000 trusted network security platform*, product datasheet, Cloud Shield Technologies Inc., Sunnyvale, California, United States.

Committee on National Security Systems (CNSS) 2012, *National Information Assurance Policy on the use of public standard for the secure sharing among National Security System*, policy, October, vol. 15.

Gasarch, W 2004, 'A survey on private information retrieval', *Bulletin of the EATCS*, January, vol. 82, pp. 72-107.

General Dynamics n.d., *TACLANE-1G (KG-175G) Encryptor*, product datasheet, General Dynamics Corporation, Scottsdale, Arizona, United States.

Global Platform Device Technology, 2011, 'TEE system architecture version 1.0', technical report, GPD_SPE_009, December.

Hutchins, EM, Cloppert, MJ & Amin, RM 2011, 'Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains', *Leading Issues in Information Warfare and Security Research*, vol. 1, p. 80-93.

King, D, Orlando, G & Kohler, J 2012, 'A case for trusted sensors: encryptors with deep packet inspection capabilities', *Proceedings of the 2012 Military Communications conference, 2012-MILCOM*, Orlando, Florida, United States, pp. 1-6.

Kushilevitz, E & Ostrovsky, R 1997, 'Replication is not needed: single database, computationally-private information retrieval', *Proceedings of the 38th annual symposium on foundations of computer science*, IEEE, Washington, D.C., United States, pp. 364-73.

Li, X, Liang, X, Lu, R, Shen, X, Lin, X & Zhu, H 2012, 'Securing smart grid: cyber attacks, countermeasures, and challenges', *IEEE Communications Magazine*, vol. 50, no. 8, pp. 38-45.

Lincoln, P, Porras, P & Shmatikov, V 2004, 'Privacy-preserving sharing and correction of security alerts', *Proceedings of the 13th conference on USENIX security symposium*, vol. 13, USENIX Association.

Ma, J, Chen, X-z & Li, J-h 2010, 'An approach to privacy-preserving alert correlation and analysis', *Proceedings of the 2010 IEEE Asia-Pacific Services Computing Conference*, (APSCC), 2010, pp. 620-24.

Madnick, SE & Donovan, JJ 1973, 'Application and analysis of the virtual machine approach to information system security and isolation', *Proceedings of the workshop on virtual computer systems*, ACM, pp. 210-24.

McKeen, F, Alexandrovich, I, Berenson, A, Rozas, CV, Shafi, H, Shanghogue, V & Savagaonkar, UR 2013, 'Innovative instructions and software model for isolated execution', *Proceedings of the 2nd international workshop on Hardware and Architectural Support for Security and Privacy*, HASP, Tel-Aviv, Israel, pp. 10:1-10:8.

Meushaw, R & Simard, D 2000, 'Nettop—commercial technology in high assurance applications', *Tech trends notes*, National Security Agency.

Niksefat, S, Sadeghiyan, B, Mohassel & Sadeghian, S 2013, 'ZIDS: a privacy-preserving intrusion detection system using secure two-party computation protocols', *The Computer Journal* March, n.p.

NIST 2001, 'Federal Information Processing Standard 140-2', Security Requirements For Cryptographic Modules, NIST, May.

Ostrovksy, F & Skeith, WE 2007, 'A survey of single-database private information retrieval: techniques and applications', *Proceedings of the 10th international conference on practice and theory in public key cryptography',* Springer-Verlag, Berlin, Heidelberg, Germany, pp. 393-411.

Scarfone, K & Mell, P 2009, 'Guide to intrusion detection and prevention systems (idps)', *National Institute of Standards and Technology special publication 800-94*, NIST, Gaithersburg, Maryland, United States.

Wang, S, Ding, X, Deng, RH & Bao, F 2006, 'Private information retrieval using trusted hardware', *Lecture notes in computer science*, ESORICS, Springer-Verlag, Berlin Heidelberg,Germany, pp. 49-64.

Williams, P, Shimeall, T & Dunlevy, C 2002, 'Intelligence analysis for Internet security', *Contemporary Security Policy*, vol. 23, no. 2, pp. 1-38.

Yasuda, M, Shimoyama, T, Kogure, J, Yokoyama, K & Kosiba, T 2013, 'Secure pattern matching using somewhat homomorphic encryption', *Proceedings of the 2013 workshop on cloud computer security*, ACM, Berlin, Heidelberg, Germany, pp. 65-76.