Volume 14, Issue2, April, 2015

ISSN 1445-3312 (Printed Journal)

ISSN 1445-3347 (On Line Journal)

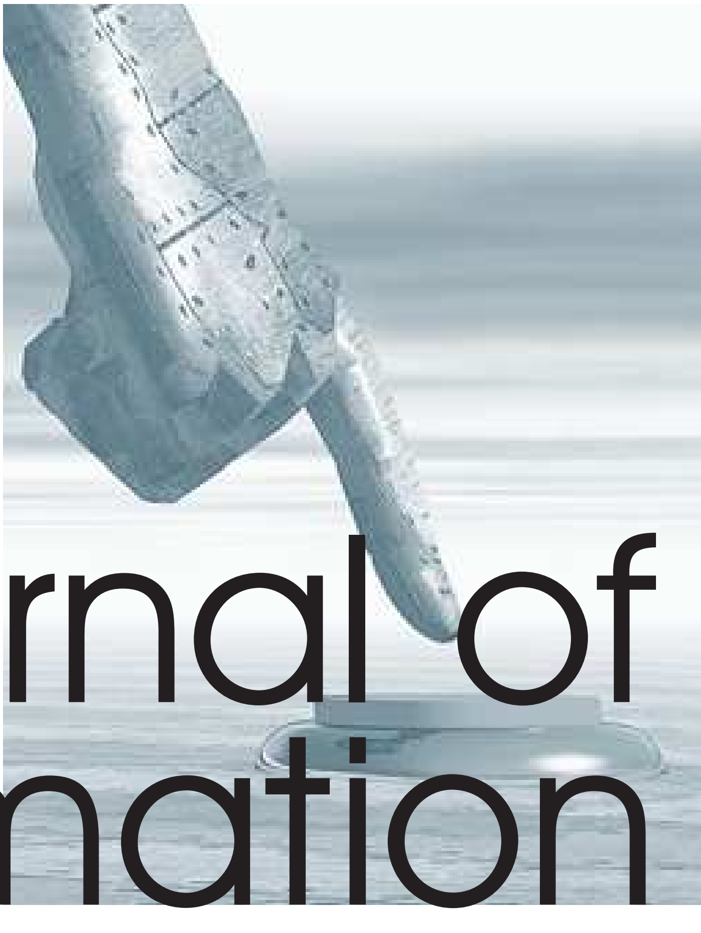

Journal of Information Warfare

These attacks can range from denial of
service to the deliberate
Disrupt of destroy data. This is similar to the
above, but disruption can be caused to the
system collecting and storing the data or
that part of the system, which disseminates
Destruction of the data can occur by physical
destruction of the storage medium, of the
data itself, so it becomes irrecoverable
to make it useless

Journal of Information Warfare (JIW)

www.Jinfowar.com

Journal Staff

Editorial Board

Advisory Board

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.

Copyright of the article is retained by the authors who warrant that they are the copyright owner and have in no way infringed any third-party copyright. In submitting the article for publication, the above warrant is implied as is the grant of a non-exclusive copyright license by the author to the *Journal of Information Warfare* to publish the work as determined by the Editorial Board.

The views expressed in contributions to this journal do not necessarily represent those of the editors, advisory board, or the publishers.

Subscription

Individual; 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 14, Issue 2

Contents

ISSN 1445-3312 print/ISSN 1445-3347 online

April 2015

In April 2014, Peregrine collaborated on a special edition of the Journal of Information Warfare (JIW), with every article being written by a serving member of the National Security Agency (NSA) staff. Our staff worked closely with the Information Assurance Directorate (IAD) personnel during a six month period to bring together an exciting issue. Here is a quote from Neal Ziring, the Technical Director for the National Security Agency's Information Assurance Directorate (IAD) on that effort - "The April 2014 issue of JIW was the first time that NSA IAD worked with an academic journal to create a special issue. It was a great learning experience for some of our internal experts, and also helped raise awareness of some of our important mission challenges among academic researchers in this field."

Once again, Peregrine is collaborating with the NSA on a new special edition of the JIW, with the nine papers as shown below. Enclosed is a list of the articles and as you can see, these papers cover key areas of concern with regards to information assurance and cyber security:

- The Future of Cyber Operations and Defense (N Ziring)
- Training Cyber Forces without Warfighting (T Walcott)
- Understanding the Co-Evolution of Cyber Defenses and Attacks to Achieve Enhanced Cybersecurity (GN Willard)
- Changing the Future of Cyber-Situational Awareness (N Newmeyer)
- The Need for Digital Identity in Cyberspace Operations (AR Friedman and LD Wagoner)
- Moving Big-Data Analysis from a 'Forensic Sport' to a 'Contact Sport' Using Machine Learning and Thought Diversity (AJ Ferguson and NM Evans Harris)
- On the Role of Malware Analysis for Technical Intelligence in Active Cyber Defense (R Fanelli)
- I Want My Smartphone. I Want It Now. And I Want to Connect to Everything from Anywhere…Now! (MLG Althouse)
- Defending Cyberspace with Software-Defined Networks (GH Bishop, SR Boyer, MJ Buhler, AJ Gerthoffer, and BC Larish)

Peregrine hopes that you enjoy this special edition. We certainly did in developing it. Cheers

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

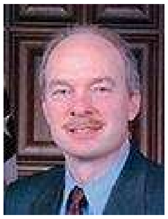

Dr. Mark L.G. Althouse is the Technical Director, Engineering, in NSA/IAD's Trusted Engineering Solutions Directorate. He is also the IAD CIA SEAM and manager of NSA's CRADA with Bechtel. Dr. Althouse graduated from Penn State University in 1981 with a B.S. in Physics. He received the M.S. in Electrical Engineering from Johns Hopkins University in 1988, and the Ph.D., also in EE, from the University of Maryland Baltimore County (UMBC) in 1995.

Garth Bishop is the lead for the National Security Agency's SDN security and policy solution. As such, Garth is responsible for shaping and implementing secure SDN practices while ensuring government security compliance. Garth is a member of the Information Technology Development Program, utilizing skill sets in vulnerability analysis and network security. Garth holds a B.Sc. in Information Technology from the University of Central Florida and expects to receive his M.E. in Electrical and Computer Engineering from Naval Postgraduate School in June 2015.

Steven Boyer serves as the lead of NSA's SDN Operational Development team. In this capacity, Steven is responsible for developing mission-critical applications for implementing SDN in a Campus Area Network as well as vital troubleshooting capabilities that will provide a seamless transition from traditional networking approaches to SDN solutions. Steven also leads NSA's OpenStack implementation of SDN where he is deploying and testing SDN solutions to explore the advantages of using SDN in an OpenStack data center environment. Steven holds a B.S.E.E from the University of Maryland.

Matthew Buhler serves as team lead for the NSA's Storage Cloud Implementation of SDN and is responsible for developing and implementing a global SDN infrastructure, providing connectivity within and between storage cloud data centers. Prior to joining NSA, Matthew worked as a systems engineer, designing test procedures and methods for network security related tasks. Matthew also worked as a network and security penetration tester, focusing primarily on network switches and routers to identify weaknesses in design. Matthew holds a B.S.E in Electrical Engineering from the University of Delaware.

Natalie M. Evans Harris is a cybersecurity leader at the NSA, with over 14 years' experience in the public sector. Currently, she is forward deployed to Capitol Hill as a Brookings Fellow, responsible for Senator Cory Booker's Cyber and Technology Innovation initiatives. Prior to this deployment, she led a tradecraft development organization responsible for creating big data analytics. During her tenure with NSA, she has been responsible for developing and deploying cyber-defensive solutions to warfighters and coordinating support between the NSA and the Department of Homeland Security. In addition to her work with NSA, Ms. Evans Harris is an adjunct professor for the University of Maryland University College, where she composes and teaches information systems security courses. She holds a Master's in Public Administration from George Washington University and two Bachelors' of Science Degrees, Computer Science and Sociology, from the University of Maryland Eastern Shore.

Robert L. Fanelli, Ph.D., is a colonel in the United States Army currently assigned to United States Cyber Command. He has served over 26 years in the Army as a signal officer and an information systems management officer and has taught computer networking and cyber security at West Point. His research interests include information system security, communications networks, and artificial intelligence.

Dr. Aaron J. Ferguson is currently serving as the Deputy Technical Director (TD) of the Fusion, Analysis, and Mitigations (FAM) Deputy Directorate at the NSA. As TD, he is responsible for providing technical leadership to the FAM leadership team, personnel, and missions, including analytics, systems and technology analysis, and operational analysis and mitigations. Dr. Ferguson holds a B.S. in Electrical Engineering from Howard University, an M.S. in Operations Research from the University of New Haven, and an M.A. and Ph.D. in Applied Mathematics and Statistics from the University of Delaware. His personal expertise areas include machine learning, software engineering, systems engineering, and risk assessments.

Mr. Arthur R. Friedman, Senior Strategist for Cyber Integration serves as the Information Assurance Directorate's (IAD) integree in the NSA Cyber Task Force (CTF), representing the interests of the IAD. He is responsible for advising CTF leadership of potential issues and ongoing activities of immediate or long-term interest. His current focus is supporting the development of information sharing strategies and policy, efforts related to cyber risk mitigation, and technologies supporting Active Cyber Defense capabilities. Mr. Friedman graduated from Hofstra University in 1979 with a Bachelor's degree majoring in Mathematics, Boston University in 1983 with a Master's degree in Business

Administration, and the Army War College in 2005 with a Master's degree in Strategic Studies.

Alex Gerthoffer serves as the lead for the NSA's campus area network SDN solution. Alex is responsible for developing and implementing SDN OpenFlow applications, which will be used to transport network traffic at the agency's branch office locations, while also focusing on increasing end-point protection. Alex is a member of the Information Technology Development Program, which has provided him with a well-rounded set of information technology skills. Alex holds a B.Sc. in Computer Security from East Stroudsburg University of Pennsylvania, and expects to receive his M. Sc. Information Assurance from Dakota State University in May 2015.

Bryan Larish serves as Technical Director for Enterprise Connectivity and Specialized IT Services at the NSA where he is responsible for setting the technical direction of the development and operation of NSA's global network infrastructure. Prior to joining NSA, Bryan worked in the Chief Engineer's office at the U.S. Navy's Space and Naval Warfare Systems Command (SPAWAR). In that role, he was responsible for implementing engineering techniques used to manage, architect, and plan the U.S. Navy's communications/IT systems portfolio. Bryan's other experience includes Technical Director for all Navy engineering policy and various engineering roles at SPAWAR. Bryan holds a Ph.D. and an M.S. in Electrical and Computer Engineering from the Georgia Institute of Technology and a B.S.E. in Electrical Engineering from Arizona State University.

Nicole A. Newmeyer is the Mission Director for the Mission Software Services Office within the Information Assurance Operations Deputy Directorate at the NSA. As Mission Director, she is responsible for guiding short term and strategic Information Assurance Operations' needs, providing technical leadership to capability development teams, and providing technical guidance to leadership. Ms. Newmeyer has held technical leadership positions in both the signals intelligence and information assurance missions at NSA. She holds a B.S. in Computer Science from the University of Maryland, Baltimore County, and is currently working toward an M.S. in Technology Intelligence at the National Intelligence University.

Dr. Larry D. Wagoner has served in a variety of technical and/or analytic organizations within the National Security Agency. Before coming to the Information Assurance Directorate (IAD) of NSA, Dr. Wagoner worked primarily in the Signals Intelligence and Research Directorates. He has a B.A. in Mathematics and Economics with specialization in Computer Science, plus a Certificate in Systems Analysis & Operations Research; an M.S. in Computer Science, as well as minors in Operations Research and Applied Mathematics; and a Ph.D. in Computer Science. All of the aforementioned degrees were earned at the University of Maryland Baltimore County (UMBC).

Thomas S. (Tom) Walcott is the Technical Director, Cyber National Mission Force, U.S. Cyber Command. He provides technical leadership and recommendations to the Commander, Cyber National Mission Force. He has a background in computer security, with experience in both intrusion detection and forensics. Dr. Walcott has a B.A. in Computer Science and Creative Writing from Beloit College; an M.S. in Computer Science from the University of California, Davis; and a Ph.D. in Computer Science from the University of California, Davis.

Gerald N. 'Chip' Willard is currently serving as a senior Technical Leader within the NSA/CSS Threat Operations Center (NTOC) supporting the Chief Operating Officer and Technical Director on strategic mission integration, technology, and research requirements. Mr. Willard is also the NSA representative to the National Science and Technology Council Committee on Homeland & National Security Interagency Working Group for Special Cyber Operations Research & Engineering (SCORE), which serves as part of the Committee's internal deliberative processes and principle venue for coordination of federal classified cyber research. Mr. Willard holds a B.S. in Information Systems Management from University of Maryland University College and an M.S. in Information and Telecommunications Systems from Johns Hopkins University.

Neal Ziring is the Technical Director for the National Security Agency's Information Assurance Directorate (IAD), serving as a 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 a technical liaison to several universities that are participants in the National Centers for Academic Excellence - Research (CAE-R) program. His personal expertise areas include router security, IPv6, VM-based secure systems, cloud computing, cross-domain information exchange, and mobile code security.

The Future of Cyber Operations and Defense

N Ziring

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

Abstract: *National and economic security of most nations have become dependent on cyberspace. Protection of cyberspace will depend, in part, on successful cyber operations. Cyberspace is the domain in which these operations take place, and it is evolving rapidly, through increased use of virtualization and cloud services, ubiquitous mobility, and the integration of cyber systems into the physical world. To be successful in this future environment, cyber operations will need to be supported by more defensible systems, to be informed by a greater understanding of system state and threat actors, and to be more adaptive.*

Keywords*:* *Cyber Operations, Cyber Defense, Future of Cyber, Situational Awareness, Defensible Systems*

Introduction

The national security and economic stability of most nations have become dependent on information systems and networks; nations rely on cyberspace for conduct of commerce, defense, intelligence, transportation, law enforcement, and many forms of social interaction (U.S. DoD 2011). In one sense, cyberspace is just another domain in which human interaction can occur, but it is not a passive background. Some elements of cyberspace are devices, processes, links, storage, and services. Safeguarding the information and interactions in cyberspace requires the creation of components and systems which can be defended, as well as responsible parties to execute these defense operations. How will the cyber environment change in the next few years, and how will cyber operations need to evolve to keep up with this change?

The U.S. Department of Defense has a succinct definition for cyber operations: “The employment of cyber capabilities where the primary purpose is to achieve objectives in or through cyberspace. Such operations include computer network operations and activities to operate and defend the Global Information Grid” (U.S. DoD 2010). This definition is very broad: it includes traditional cyber operations, such as maintaining a network's configuration and defending it from external attack, but it also includes operations that create effects outside cyberspace. Cyberspace has become an operational domain for many reasons, but two main ones are the growth of cyber threats, and the increased risk they pose to highly dependent societies. First, the scope of threats, their variety and persistence have increased dramatically just in the past few years (Symantec 2014). Secondly, greater reliance leads to greater risk, partly because a cyberattack can cause harm beyond assets resident in cyberspace. Indeed, the negative effects from a cyberattack or compromise can occur very quickly. All of these trends imply that organizations of all sizes must exercise constant watchfulness and must be prepared for immediate and informed response.

This short paper introduces several of the topics covered in this journal and tries to illustrate the context into which they fit. The future of cyber operations will be much more complex than its past. The environment in which cyber operations will be conducted, as well as the systems, networks, and services that comprise the cyber domain, will be increasingly integrated with each other and with the physical world. Virtualized substrates will become the norm, and software-defined, dynamic platforms and networks will underpin them (Knorr 2013). Security for these systems, and for the devices that depend on them, will increasingly depend on identity. Operations within cyberspace will depend on understanding the posture of the systems involved and on adaptation to the actors opposing them.

The Past, Present, and Future Cyber Environment

Modern cyber operations began in the computer network environment of the late 1990s and early 2000s, during a time when networks were relatively static and managed as individual enclaves. Cyber operations were largely concerned with keeping external threat actors out of a network (that is, preventing initial penetration). Networks of that period had stable configurations, were mostly homogeneous, and had assets that were managed by dedicated administrators. These assumptions are implicit in cyber operations practices as illustrated by the organization of the military academies' Cyber Defense Exercise in its early years (Schepens & James 2003). Each enterprise ran its own IT, and ‘mobility’ commonly meant a laptop tethered back to the enterprise network. Identities were typically issued by an enterprise and valid only within that enterprise or with particular partners.

Today's environment presents some significant differences and imposes many new challenges for cyber operations. There are several significant trends, each disruptive in its own right:

1. Widespread virtualization—most servers in enterprise datacenters are now virtualized, and many other forms of virtualization are becoming commonplace (McLellan 2013). This creates new challenges for cyber operations because assets are more dynamic, communication paths are more complex and harder to monitor, and relationships between assets are more fluid. Also, the virtualization layer adds complexity and new attack surface to operational networks.

2. Migration to cloud services—since about 2006, enterprise and commercial cloud services have gained adoption at a steady pace in industry. Since the introduction of the Federal Risk and Authorization Management Program (FedRAMP) in 2012, the U.S. federal government has also been migrating services to commercial clouds (Taylor 2014). The multi-tenant, resource-sharing nature of clouds adds many complications to cyber

operations. Technically, clouds add a new technology base, additional points of interactions, and extra layers of infrastructure that cyber-operations personnel must understand. The presence of multiple tenants complicates operations because assets and services within the scope of an operation are commingled with those that are out of scope.

3. Rise of mobility—users now expect full access to information and services regardless of physical location. The rapid growth of smartphone use, and the associated growth in mobile services, has imposed several changes on the cyber environment that affect cyber operations. First, the boundary between enterprise systems and personal systems has blurred—sensitive information does not stay confined in the enterprise where policies can be enforced, but may be copied to mobile devices and cloud services. More subtle, but no less disruptive, is the inconstant nature of mobile connectivity. Mobile devices use multiple networks, and are frequently inaccessible. From a cyber-operations viewpoint, mobile devices constitute a shifting swarm of accesses and interactions where traditional tools and techniques do not apply.

4. Growth of connected, cyber-physical systems—in the last few years, devices and systems that affect the physical world have become increasingly accessible from the Internet. While industrial control systems have been in common use for decades, for most of that time they were isolated from most threats by lack of connectivity. Furthermore, many other types of physical services are now mediated by cyber systems: building control, transportation, and various facets of critical infrastructure.

These trends will continue for the near future, and some will accelerate. Cloud computing will become ubiquitous, even for sectors where security concerns had previously discouraged adoption (Kenyon 2014). Mobility will become the norm, in the sense that every enterprise will allow at least some of its assets and information to span diverse mobile devices. These trends, together, will create intense pressure for globally valid identities and associated authentication and authorization services. Networks will become more dynamic, driven by flexible technologies such as Software-Defined Networking (SDN) and Automatically Switched Optical Networking (ASON). The growth of cyber-physical systems will also continue; increasingly, devices and services that interact with and control elements of the physical world will be accessible from and subject to attack from cyberspace. This trend, often called 'the Internet of Things', involves all manner of extension into the physical world of devices with computing, sensing, and actuation capabilities (Miorandi *et al.* 2012). The integration of cyber and physical domains will greatly increase the potential scope for cyberattacks and will impose a corresponding need for defensive operations.

Operating in Cyberspace

Operating in the cyber environment of the future will require significant evolution of current practice. Three related areas of change will be especially important. First, systems must be designed and built to support operations, particularly time-sensitive operations such as incident response. Second, cyber operations must become extensively data-driven. Operators and automated systems that support them must be enabled with extensive data analysis, and the analyses must incorporate both local and global context. Third, systems and operations must become more adaptive.

Effective cyber operations, especially defensive operations, depend on accurate and timely knowledge of the operational context (usually a target network or enterprise), and on the ability of the operational environment to support actions (such as defensive responses). A network that is designed and built to facilitate operations in its own defense, including monitoring and response, is said to be defensible.

Understanding the operational posture of a network, including its weaknesses and the status of defenders and attackers, must be based on data and science. A prescient article from 2004 identifies key attributes of scientifically grounded cyber security, and the principles in it apply directly to today's situation over a decade later (Saydjari 2004). Six core elements of cyber defense are described there; three of them are especially important for the future of cyber operations:

1. Sensors and data collection—cyber operations must be supported with accurate information about the particular networks being defended, but also about the global context. Sensors emplaced in a defensible network are essential for the local view. For the global context, it will be necessary to fuse multiple sources: commercial reputation data, threat intelligence, and partner security posture.

2. Situational awareness—collected data must be analyzed to produce coherent and actionable information to support operational decisions. Today's system logs, sensors, and intelligence sources provide large volumes of data; a key aspect of analysis is to filter out noise and to present the most relevant results to operators. Big data analytic technology can provide an effective platform for situational awareness (Roddy 2014).

3. Defense mechanisms—accurate situational awareness informs action, but cyber operators need effective and reliable means to execute actions. A defensible network includes specifically designed and deployed mechanisms for controlling the assets that comprise the network, as well as orchestration for applying multiple mechanisms in concert.

Protecting information is a critical aspect of defense for most networks. System designers and cyber operators must explicitly consider safeguarding of information, separately from the systems which host it. There are three main elements to safeguarding information: designating the protection required for information objects, designating entities and their rights to access information, and enforcing the access policies applicable to those objects and entities. There are many strategies for this, but one which has proven highly effective at NSA is Attribute-Based

Access Control (ABAC). In an ABAC model, information objects bear simple tags, and entities are assigned attributes based on their rights and authorities. Policies express which attributes entities must possess to gain access to information with certain tags (Sandhu, Ferraiolo, & Kuhn 2000).

Finally, operations must be adaptive. Threat actors adapt to defenses, for example, using obfuscation to evade anti-virus software or adapting denial-of-service tactics to defensive measures (Engelman 2012). Defensive operations must be similarly flexible.

It is essential that defenders analyze new tradecraft (including malware) and be ready to adjust and combine defensive measures to detect and defeat it. Adaptation must be supported at several levels: from basic network operations, up through big data analytics and intelligence analysis. Fortunately, several key technologies already exist to help make defensible networks and cyber-operations tradecraft more adaptable. Software-Defined Networks will allow defenders to adapt network topologies to enforce new policy, to douse undesirable dataflows, or even to randomly change networks to defeat attackers (Jafarian, Al-Shaer & Duan 2012). Also, machine learning will allow analytics to spot anomalies in time to enable response before attackers achieve their objectives. These and other technologies can support highly-adaptive operations. Then, the challenge will be training cyber-operations personnel to use these technologies effectively.

The Future of Cyber Operations and Achieving Long-Term Security

Dan Geer (2014), noted security authority, has defined a state of security as the state where there is “an absence of unmitigatable surprise”. The systems that comprise the cyber environment will always be subject to attack and to compromise. This does not excuse system designers from striving to create secure, defensible systems, nor does it reduce system operators' obligations to maintain secure configurations. There are many excellent published works on how to accomplish these, such as the *Community gold standard framework* (NSA 2014). Secure design and configuration can and do defeat attackers and raise their costs; plus they provide critical capabilities to enable cyber operations. Nevertheless, operations will still be necessary. Using Geer's definition, no system can be made absolutely immune to surprise, but effective operations can mitigate the remaining surprises.

Achieving security in practice, and on large scales, will be exceptionally challenging. The future cyber environment will be highly diverse and dynamic, and its integration with large swathes of society and with the physical world will expand and amplify potential impacts of attacks. There will be three essential elements to security in that environment:

1. secure and defensible systems;

2. timely and sustained understanding of both the environment being secured, and the threat actors who may attack it; and

3. effective and adaptive cyber-operations personnel and supporting tools.

Fortunately, the community possesses all the right building blocks for these essential elements. The papers in this issue of the *Journal of Information Warfare* present detailed analyses and experience with many of them. With attention to the continued development and application of these elements, both the public and private sectors can reap the benefits of cyberspace while maintaining social, economic, and national security.

References

Engelman, E 2012, ‘Bank cyber attacks enter fifth week as hackers adapt to defenses’, *Bloomberg*, 19 October 2012, viewed 1 March 2015, <http://www.bloomberg.com/news/articles/2012-10-18/bank-cyber-attacks-enter-fifth-week-as-hackers-adapt-to-defenses>.

Geer, D 2014, ‘APT in a world of rising interdependence’, speech delivered at NSA 26 March 2014, Odenton, Maryland, United States, viewed 1 March 2015, <http://geer.tinho.net/geer.nsa.26iii14.txt>.

Jafarian, JH, Al-Shaer, E & Duan Q 2012, ‘Openflow random host mutation: transparent moving target defense using software defined networking’, *Proceedings of the first workshop on Hot topics in software defined networks,* ACM, 2012. pp. 127-32.

Kenyon, H 2014, ‘DoD changes cloud computing policy’, *InformationWeek Government*, 17 November 2014, viewed 1 March 2015, <http://www.informationweek.com/government/cloud-computing/dod-changes-cloud-computing-policy/d/d-id/1317511>.

Knorr, E 2013, ‘Ahead of the curve: 9 trends for 2014 and beyond’, *InfoWorld*, 4 November 2013, viewed 1 March 2014, <http://www.infoworld.com/article/2612875/cloud-computing/9-trends-for-2014-and-beyond.html>.

McLellan, C 2013, ‘Virtualizing the enterprise: an overview’, *ZDNet*, 1 August 2013, viewed 1 March 2015, <http://www.zdnet.com/article/virtualizing-the-enterprise-an-overview/>.

Miorandi, D, Sicari, S, De Pellegrini, F & Chlamtac, I 2012, ‘Internet of things: vision, applications and research challenges’, *Ad Hoc Networks*, vol. 10, no. 7, pp. 1497-516.

National Security Association 2014, *The community gold standard framework*, version 2.0, Information Assurance Directorate, National Security Agency, Fort Meade, Maryland, USA, June 2014, viewed 1 March 2015, <https://www.iad.gov/iad/CGS/cgs.cfm>.

Roddy, S 2014, 'How IAD leverages big data for anomaly and malware detection', *Journal of Information Warfare*, vol. 13, no. 2, April 2014, pp. 70-75.

Sandhu, R, Ferraiolo, D & Kuhn, R 2000, ‘The NIST model for role-based access control: towards a unified standard’, *ACM workshop on role-based access control*, vol. 2000, n.p.

Saydjari, OS 2004, ‘Cyber defense: art to science’, *Communications of the ACM*, vol. 47, no. 3, March 2004, pp. 52-57.

Schepens, WJ & James, JR. 2003, 'Architecture of a cyber defense competition', *IEEE International Conference on systems, man and cybernetics, 2003.*, vol. 5, IEEE, 2003.

Symantec Corporation 2014, *Symantec internet security threat report 2014*. Symantec Corporation, vol. 19, April 2014, p. 13, viewed 1 March 2015, <http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf>.

Taylor, L 2014, ‘FedRAMP: history and future direction’, *IEEE Cloud Computing* vol. 3, pp. 10-14.

United States Department of Defense 2010, Chairman of the Joint Chiefs of Staff, *Joint terminology for cyberspace operations*, viewed 1 February. 2015, <https://info.publicintelligence.net/DoD-JointCyberTerms.pdf>.

——2011, *Department of Defense strategy for operating in cyberspace*, U.S. Department of Defense, July 2011, viewed 1 February 2015, <http://www.defense.gov/news/d20110714cyber.pdf>.

Training Cyber Forces without Warfighting

T Walcott

Cyber National Mission Force
United States Cyber Command
E-mail: JIWfeedback@nsa.gov

Abstract: *Collective and individual training for military cyber operations poses challenges not faced by industry, academia, or other governmental areas. The warfighting mission comes with unique issues scarcely dealt with by the modern United States, such as foreign attacks taking place on United States' infrastructure. As a result, there are limited existing processes to draw upon. Effective training is further hampered by lack of operational experience. This paper discusses the challenges of gaining experience in cyber operations, explores several avenues for obtaining real-world operational experience outside of warfare, and considers the applicability of those operational scenarios to training.*

Keywords: *Cyber Operations Training, Cyber Forces Training, Cyber Operations other than War*

Introduction

Standing up the United States' Cyber Command requires identifying, recruiting, educating, and training military personnel to operate in cyberspace. The challenges of identifying, recruiting, and educating talented individuals are not unique to the military and have received attention from outside of the military community. Considerable effort throughout both the public and private sector has been directed towards identifying and recruiting individuals, and there are a variety of educational resources available to improve individual proficiency. (For a more thorough discussion of the evolving educational requirements and resources, readers may see, for example, Kallberg & Thuraisingham 2012.)

During the battle for Monterrey, Mexico, in 1846, the United States' Army found itself in urban warfare. The U.S. officer corps was decimated over the course of one day of fighting in this novel environment. Yet Texas volunteers participating in that very battle were familiar with fighting in this environment; the tactics they recommended, once incorporated into the collective functioning of the forces, helped turn the tide (Dishman 2010). The cyber domain is no different. However, learning from bitter experience, while effective, is not optimal.

The Need for Training

> Train like you fight and fight like you train. (Military adage)

Building a capable military force in a new domain requires both education and training. There is a critical distinction between the two: education is learning why to do something; training is learning how to do it. There are significant external resources available for and devoted to education—the academic system and conferences, for instance—with robust cyber curricula. The high demand for computer security experts and high visibility for security lapses are industry incentives to continue growing education resources.

Unfortunately, those educational environments do not immediately translate to the training needs of our forces. This holds true in both individual training, designed to engage academic skills in an operational context, and collective training. The military places great emphasis on both types of training, but it is the collective training that distinguishes a military unit from an armed mob. An analogy in cyber operations might be contrasting a Network Operations Center (NOC) and Anonymous.

Briefly, an NOC has the responsibility for maintaining situational awareness of and responding to network incidents. A group of individuals of varying skill levels will dynamically partition (and re-partition) tasks based on urgency and complexity in order to maintain network health. Because an NOC is generally size constrained, each team member must be capable of effective contribution; a failure of one team member may negatively impact the entire team.

Anonymous is an extreme example of decentralization, sacrificing structure for raw numbers of participants. Activities will be advertised at the whim of any member, have no guarantee of support from other members, and have no particular oversight. Of note, the Anonymous structure is much more forgiving of less trained participants. While individuals can boost the achievements of the whole, an individual failure has very limited impact on the overall effort.

The U.S. military is designed to operate in a structured fashion. Military operations are predicated on a chain of command, rules of engagement, areas of responsibility, and a clear division of effort. This structure helps ensure coherent execution of mission, compliance with legal obligations, and accountability. Thus, applying military structure to collective training within the cyber domain is of paramount importance for successful employment of cyber forces. This application of structure to collective training efforts also means that failures anywhere within the team may damage overall team performance.

An interesting possibility is a hybrid approach, with a small cadre of experts trained in this particular discipline who present specific capabilities not represented by other team members. There is existing precedent for this possibility with most sophisticated modern munitions; the individuals responsible for deploying a missile or bomb are generally not qualified to build them. The cyber domain does present some unique risks to this division of expertise. Capabilities in cyberspace are extremely perishable; there are few, if any, kinetic environments where firing one missile can result in invalidating all future missiles of a given design. Managing that risk requires a broad suite of specialized skills throughout design, planning, and execution of cyberspace missions. While a hybrid approach might leverage success in individual training, it would require significant rethinking of how to compose and employ the military cyber-forces from the standpoint of practical implementation. Therefore, such a construct is unlikely to be viable in the near-term.

As previously mentioned, legal concerns are a critical component of both education and training. While a detailed discussion is beyond the scope of this paper, legal analysis of cyberspace operations are ongoing and a rich area of both discussion and publication. Interested readers may wish to review Lin (2010) and Schmitt (2013).

The Challenges with Training

The cyberspace domain is unlike any other. For instance, the effective lifetime of most cyber capabilities is much shorter than effective lifetimes of traditional capabilities. It is hard to imagine a traditional munition which the firing of runs the risk of generating global immunity to that munition. Certainly, the traditional military acquisition process would not be entirely comfortable with a model that required exercising after every operation; the acquisition programs for current platforms, such as naval vessels or aircraft, are not known for being nimble and may take years to establish. Yet in cyberspace, new tools and techniques are regularly mitigated in hours or days. While this state of affairs has implications that resonate throughout many military processes, it certainly also requires specialized training for those forces responsible for defending against and employing cyber capabilities.

According to one clause of the U.S. Cyber Command mission statement, USCYBERCOM “conduct[s] full spectrum military cyberspace operations in order to enable actions in all domains”. Such a task represents a unique national responsibility that covers a broad range of missions. The complexity of this task is compounded by the fact that “most aspects of joint operations rely in part on cyberspace” (U.S. DoD 2013). Just as a commercial jet pilot is not trained to fire or evade missiles, so commercial security practices are not necessarily translatable across the entire environment where USCYBERCOM will operate. As a result, those responsible for developing training regimes have a less-than-complete understanding of the full suite of skills required to operate in the cyber environment.

Another challenge with training is that the value of any training drops when it does not reflect real-world conditions. It is difficult even to evaluate a proposed training regimen or environment in the absence of relevant real-world experience. While this reality is a complication for individual training, its effects on collective training are compounded. Each variance between theory and practice at the individual level leads to accumulated differences between team elements.

One option is to train forces with the best knowledge available to date, and hope that they are called upon to fulfill a similar real-world mission in the context of warfighting. This option is unlikely to result in an ideal outcome. The decision to employ cyber forces in wartime will represent a significant policy decision that, on the balance, will only become acceptable under exigent circumstances. It is unlikely that exigencies will be neatly aligned with past training scenarios, and high-stakes circumstances are generally not the optimal time for the first operational test of military capability.

Some of this risk could be mitigated by training for exigent circumstances. This approach to training, however, results in a force with a focus that (by definition) is only appropriate *in extremis*. Such scenarios are the bread and butter of major training exercises, and rightly so as they stress-test the seams of collective training. Yet the bulk of any warfighter's time is not spent firing at an adversary; most of the time is spent optimizing the chances of success. In that sense, a training scenario focusing on exigent circumstances does not prepare the warfighter well for his or her normal day-to-day routine.

Another training option is to learn from cyber-warfighting experience. Although this type of training is critical for refining training and improving forces, it is an undesirable way to start learning. There are issues with committing a force to battle that is (by definition) untrained. It would also mean requesting a significant policy decision that would perforce be executed by a poorly-trained force; the political calculus shows very clear risks and very abstract benefits.

The U.S. Army's first encounter with urban warfare was the Battle of Monterrey on 21 September 1846 against Mexico. Major Luther Giddings recounted his experience in the 14 November 1846 issue of the *Niles Register*:

> We moved rapidly through a labyrinth of lanes and gardens, without knowing or seeing upon what point of the enemy's line we were about to strike. At every step, discharges from the batteries in front became more deadly.

Losses were substantial until recommendations from experienced Texas volunteers were implemented on 23 September. These included one team's breaking through building walls while another deterred response through suppressive fire; the entry team would then seize the building roof. This strategy permitted relatively safe vantage to surround and to target a garrison, an action that "would have made bloody work of [the Americans] had [they] used the tactics of the 21st" (Dishman 2010, p. 180). Some of the tactics learned in the Battle of Monterrey are still employed today in urban warfighting.

Cyber warfare is analogous to 19th-century urban warfare; each presents a new domain where operational military experience is limited. Yet the costs of failure can be considerably higher in cyber warfare, given the "critical dependence on cyberspace, for the US in general and the joint force in particular" (U.S. DoD 2013).

The situation appears to be a catch-22 where forces cannot optimally train without warfighting experience and cannot gain warfighting experience absent training. This paper seeks to break this impasse by identifying opportunities to gain and refine experience in comparatively low-risk environments. While such experience would not perfectly mirror warfighting, it is a significant improvement over a first engagement without any experience at all.

Building Experience outside of Warfare

> The United States employs its military capabilities at home and abroad in support of its national security goals in a variety of operations that vary in size, purpose, and combat intensity. The use of joint capabilities in **military engagement, security cooperation, and deterrence** activities helps shape the operational environment and keeps the day-to-day tensions between nations or groups below the threshold of armed conflict while maintaining US global influence. […] The associated general strategic and operational objectives are to **protect** US interests and **prevent** surprise attack or further conflict. (U.S. DoD 2010)

As is highlighted in the excerpt above, the United States' military conducts a variety of activities intended to prevent armed conflict—that is, by definition, there are activities that are not warfare, but fall within the domain of military operations. This should be no great surprise; the amount of time that military members spend literally fighting an adversary is trivial. Most of the time is spent in preparation, be that acquisition, maintenance, logistics, training, intelligence, security measures, or any number of other activities undertaken by military personnel.

There are a set of non-warfare activities that are operational in nature. These are outlined in U.S. DoD Joint Publication 3-07, *Joint doctrine for military operations other than war*. This publication highlights several opportunities for operations that might chip away at the uncertainty and might help resolve the experience/training deadlock in the cyber domain. Exploring such avenues may not perfectly capture the warfighting environment, but can provide real-world operational and training scenarios that improve force readiness and capacity and also better model the tasks on which forces will spend the bulk of their time.

To that end, below are a few selected examples of military operations other than war. For each, there is a representative excerpt of the definition, some discussion as to how those operations might translate into the cyber domain, and an assessment of whether such operations might be suitable for cyber forces. The discussion is confined to operational relevance and implementation and presumes appropriate authorities.

Humanitarian assistance

Humanitarian Assistance (HA). HA operations relieve or reduce the results of natural or manmade disasters or other endemic conditions such as human pain, disease, hunger, or privation is countries or regions outside the United States. […] DOD provides assistance when the relief need is gravely urgent and when the humanitarian emergency dwarfs the ability of normal relief agencies to respond. (U.S. DoD 1995)

Joint Publication 3-07 (U.S. DoD 1995) also notes that “US forces can provide logistics; command, control, communications, and computers”. Given the increasingly significant role of the Internet as a core communications infrastructure, finite bandwidth limitations, and the taxation of that finite bandwidth in disasters, there may be times during which USCYBERCOM forces could manage bandwidth and network security to ensure humanitarian efforts can communicate when resources are scarce.

Disruptions in the cyber domain contribute to humanitarian crises. Modern logistics are highly dependent upon reliable networks, whether they are electrical, or are routed through the telephone or the Internet. Providing rapid resumption of services—which could include security assurances for networks or additional options for routing data—can contribute significantly to easing a crisis. For the USCYBERCOM forces, operating in a degraded environment provides experience in resilience and adaptability often valuable in traditional armed conflict.

There are similarities between the technical support requested in humanitarian relief efforts and in network incident response. In both cases, military responders must simultaneously work with existing system administrators and others who are familiar with local network usage and critical requirements, and operate under field conditions. That local knowledge critically supplements the specialized expertise needed to mitigate network threats. While the threats in each scenario are quite different, the time pressures and operational flexibility necessary are similar.

Finally, the presence of dedicated information security personnel will—from a technical perspective—impose no additional risk. This type of support could, at the very least, provide some measure of risk management. For these reasons, support to humanitarian assistance may provide valuable operational experience and inform future training for cyber forces.

Shows of force

> Show of Force Operations […] operations, designed to demonstrate US resolve, involve increased visibility of US deployed forces in an attempt to defuse a specific situation that if allowed to continue may be detrimental to US interests or national objectives. (U.S. DoD 1995)

There is no doctrinal definition for a show of force in cyberspace. Still, a show-of-force cyberspace operation could be defined nearly identically to the definition of traditional show-of-force operations, with the substitution of “cyberspace forces” instead of (or in addition to) “deployed forces”. One challenge that is unique to the cyber domain is ensuring unambiguous attribution; this is generally not a problem with aircraft carriers, but is an issue with almost every network activity. More formally, one characteristic of show-of-force operations should be attribution and non-repudiation.

There are a variety of risks included in an aggressive show of force, such as misperceptions of an exercise as genuine activity. In the cyber domain, these challenges may be exacerbated by still-nascent international norms and processes. There is no well-seasoned set of best practices, and the decision cycle for online activity is short. Briefly, if other entities are in a nascent operational and training state, a show of force could provide pressure that catalyzes undesirable action.

A show of force is not necessarily aggressive. Post September 11, 2001, there was an enhanced security posture at airports. That activity certainly might have had a deterrent effect on would-be terrorists, but it also may have contributed to defusing a tension felt by the general public. Whether or not that was an explicit objective, that show of force served a defensive purpose and acted as a very visible deterrent.

If, for example, public anonymizers used to connect to Department of Defense web sites delivered both the web page and a banner indicating that the use of anonymous browsing techniques was generally subject to increased logging activity, it is unlikely that the statement would be perceived as warfighting *per se*. Such a page and banner, however, would convey that the Department of Defense could identify the relevant behavior.

Shows of force have the advantage of taking place outside of a classroom environment. Carrying out a low-risk operation in public view amidst civilian activity is a capability both difficult to model and valuable to cultivate. There is some level of political risk that accrues to publicly visible activity, yet the risk is significantly less than that posed by equivalent operations in wartime.

Enforcement of sanctions

> Enforcement of Sanctions […] operations […] employ coercive measures to interdict the movement of certain types of designated items into or out of a nation or specified area. (U.S. DoD 1995)

The first challenge in this discussion is interpretation of the definition in the context of cyberspace. The language of sanctions enforcement is specifically limited to a border crossing (“into or out of” an area), which creates certain inherent limitations. Could types of information be designated items?

The legal questions underlying this issue have been outlined in Benatar and Gombeer (2011), but remembering the limitations of technical possibilities informs the legal discussion with a reminder of constraints in implementation. From a technical standpoint, a perfect solution cannot be developed for border interdiction. Interdicting data in cyberspace is extremely difficult while the data is in transit; identifying data as sanctioned might take significantly longer than the transfer of the data to its destination. Putting data in quarantine for the duration of examination may help, but will not be a perfect solution. Image files with embedded steganography provide a simple example. Running some set of detection algorithms against all image files is computationally prohibitive. The outcome would likely be unacceptably long quarantine periods.

A sanctioning body may, of course, determine that the heightened cost imposed by sophisticated concealment of sanctioned information is an adequate penalty. After all, no real world interdiction process is perfect. The object is to impose an unacceptable (or unsustainable) cost to an activity. Given that end, cyberspace sanctions could complement traditional interdictions.

Sanctions enforcement would therefore be a useful operational activity, given the lack of viable alternatives in cyberspace and the acknowledgement that no sanctions regime can be perfectly enforced. Such an operational activity requires a clear understanding of the technical limitations to the approach and the tradeoffs that must be made in implementation. This would provide a useful example for policymakers of how capabilities in cyberspace may supplement more traditional, well-understood capabilities. Lessons learned in sanctions enforcement could also have implications for defensive cyber operations.

Enforcing exclusion zones

> An exclusion zone is established by a sanctioning body to prohibit specified activities in a specific geographic area. Exclusion zones can be established in the [air, sea, or on land]. The purpose may be to persuade nations or groups to modify their behavior to meet the desires of the sanctioning body or face continued imposition of sanctions, or use . . . [of] threat of force. (U.S. DoD 1995)

The above definition of exclusion zones does not include cyberspace, but it could be extended to do so—at least in a legislative sense. Such an extension would highlight a major technical challenge.

One commonly used exclusion zone is an air exclusion area. The U.S. military enforced an air exclusion area for Iraqi forces when Saddam Hussein was in power, for example. Such a wide-ranging activity would not be sustainable in the cyber domain. There are too many options for accessing cyberspace. Wired connections can be disrupted with relative ease, but wireless connections and satellite connections are substantially more challenging. For more granular tasks, significant care must be taken to avoid running afoul of the issues seen in sanctions enforcement. By way of example, reliably identifying a single user's network traffic can be quite difficult; doing so in real time is impractical at scale.

There are some more encouraging examples that make exclusion zones worth examining. The Stuxnet virus provides an example of how an enrichment activity might be prohibited, and also some worthwhile cautions. The Stuxnet virus was engineered to seek out and to alter the behavior of Siemens S7-300 SCADA systems with variable-frequency drives from vendors Vacon or Fararo Paya when those drives were operating in a specific frequency range (Chien 2010). That level of precision in targeting demonstrates a technical capacity to be extremely discriminating in behavior and activity—except, of course, that Stuxnet had no cognizance of geography.

There are several challenges to conducting such an operation—not the least of which is identifying where excluded materials reside in both a network and a geographic sense, or identifying unique characteristics of the excluded materials that permit automated identification and response. Stuxnet once again proves illustrative, as it spread beyond the presumed initial infection point in Iran.

Conducting a similar activity as an enforcement of exclusion zones would encourage a robust public discussion—and would increase the cost of the sanctioned entity's operation, as it would be compelled to adopt a heightened security posture for a presumably highly valued initiative. While this might not fully interdict the excluded activity or material, the imposition of additional costs for acquisition and storage might be desirable in and of itself (as was the case with enforcement of sanctions).

From a force-development standpoint, the enforcement of exclusion zones is a comparatively complex operation. It requires discriminating processes and procedures. These factors make exclusion zones a poor choice for an initial operational example—but a fine study for sophisticated operations.

Conclusion

Effective individual and collective training is foundational for a skilled military. Absent operational experience, assessing the realism and effectiveness of current training for cyber forces is difficult.

Looking at four examples of military operations that are other than war effectively highlights opportunities for cyber operations outside of the warfighting domain. Applying these operational concepts to the cyber domain will require careful thought, because of the dearth of experience. Gaining experience, however, is a crucial step to ensuring that training reflects real-world operations and to maximizing the chances of success in any cyberspace-based military engagement.

References

Benatar, M & Gombeer, K 2011, 'Cyber sanctions: exploring a blind spot in the current legal debate*', 4th Research Forum of the European Society of International Law*, viewed 9 February 2015, < http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1989786>.

Chien, E 2010, *Stuxnet: a breakthrough*, Symantec blog, 12 November, viewed 2 December 2014, <http://www.symantec.com/connect/blogs/stuxnet-breakthrough>.

Dishman, CD 2010*, A perfect Gibraltar*, University of Oklahoma Press, Norman, OK.

Giddings, L 1846, 'Sketches', 14 November, *Niles Register.*

Kallberg, J & Thuraisingham, B 2012, 'Towards cyber operations—the new role of academic cybersecurity research and education*', Proceedings from 2012 IEEE International Conference on Intelligence and Security Informatics* (ISI 2012), viewed 9 February 2015, <http://works.bepress.com/jan_kallberg/9/>.

Lin, H, 2010, 'Offensive cyber operations and the use of force', *Journal of National Security Law & Policy*, viewed 9 February 2015, <http://jnslp.com/wp-content/uploads/2010/08/06_Lin.pdf>.

Schmitt, MN, ed. 2013, *Tallinn manual on the international law applicable to cyber warfare*, Cambridge University Press, Cambridge.

United States Department of Defense, Joint Publication 3-0, 22 March 2010, *Joint operations*, United States Department of Defense.

——, Joint Publication 3-07, 16 June 1995, *Joint doctrine for military operations other than war*, United States Department of Defense.

——, Joint Publication 3-12 (R), 5 February. 2013, *Cyberspace operations*, United States Department of Defense.

Understanding the Co-Evolution of Cyber Defenses and Attacks to Achieve Enhanced Cybersecurity

GN Willard

NSA/CSS Threat Operations Center
National Security Agency, Fort Meade, Maryland, United States
JIWfeedback@nsa.gov

Abstract: *This article examines the notion of cyberattack-and-defend co-evolution as a mechanism to better understand the influences that the opposing forces have on each other. The concept of co-evolution has been most commonly applied to a biological context involving living organisms and nature-based adaptations, but it can be applied to technological domains as well. Cybersecurity strategies based solely on technological features of attack-and-defend adaptations do not immediately reveal a co-evolutionary relationship and are typically seen more as cyber arms races. In order to leverage cyber co-evolution in support of cybersecurity, the human-driven behaviors of cyberattack-and-defend adaptations have to be incorporated. In other words, the mission must serve to drive human motives and goals, and in many cases, must limit the scope of an attacker's adaptations.*

Keywords: *Cybersecurity, Co-Evolution, Moving Target, Game Theory, Cyber Deception, Information Operations*

Introduction

A cyberattack is defined as any malicious act targeting a network's confidentiality, integrity, or availability. Various groups within the cybersecurity community make a distinction between Computer Network Exploitation (CNE), most commonly associated with information theft, and Computer Network Attacks (CNA), typically associated with disruption or destruction of information systems; however, for the purposes of this article, these distinctions are not essential. For convenience, therefore, these two forms of malicious activities will be abstractly and collectively grouped as cyberattacks. In the early days of the Internet, most of the initial computer-based attacks were in the form of computer worms and viruses primarily intended to gain notoriety for the malicious code author and to expose vulnerabilities of popular software and hardware makers in order to embarrass them. While these attacks could result in lost data, service interruptions, and lost productivity, they were mostly seen as simple acts of vandalism and great annoyance. These kinds of problems made the job of system security more of a technological problem that focused on vulnerabilities rather than on threats, which eventually led to the now endless cycle of software and firmware updates and patches. Now that there are more purposeful attacks by criminal, terrorist, and state-sponsored threat actors, the human motives and goals of malicious cyber behavior must be considered in formulating cybersecurity strategies.

Mounting anecdotal evidence indicates that malicious cyber actors learn, adapt, or, in other words, react to the defensive measures put into place by the cybersecurity community as much as network defenders react to attacks. When coupled with the human behavior factor, the constant cyclic attack-defend-attack behavior reveals a demonstrable co-evolutionary relationship between cyber-attack and defense-development activities. This article proposes to take advantage of the attack-defend co-evolution phenomenon by focusing on an understanding of the attacker's response to defensive measures in the context of the attacker’s mission goals and objectives. Developing an understanding of these missions and goals will generate greater predictive analysis capabilities and, more importantly, better means to influence the attacker’s evolution in a manner that plays into cyber-defensive strengths.

Rules (or Laws) of Cybersecurity

In order to have a thoughtful discussion about cyber co-evolution, it is important to explicitly lay out some innate rules of cybersecurity, which will have relevance throughout this article. Recognition of these axioms should help to remove some of the self-imposed constraints that have perhaps limited the thinking and the progress of the cybersecurity community.

Rule #1: They are going to get in.

The focus of cybersecurity has long been on keeping malicious code and hackers from gaining access to systems. This focus has led to a strategy of developing defensive (mostly detection) capabilities at the perimeter of networks, particularly gateway connections to external networks. Often there is a network ‘de-militarized zone’ (DMZ) between an internal business network and the (public) Internet. These network defense strategies have further evolved in some cases by creating closed, special-purpose network environments (almost) inside the corporate network to provide additional layers between an organization’s most precious assets and the public network. Still, malicious stuff continues to get into networks.

The threats to networks have become more organized, more sophisticated, and better resourced. In fact, these threats have been given a name: Advanced Persistent Threats (APTs). APTs are cyberattacks mounted by organizational teams that have deep resources, advanced penetration skills, specific target profiles, and remarkably persistent efforts. These threats tend to use sophisticated custom malware that can circumvent most defenses and stealthy tactics, as well as demonstrate good situational awareness by evaluating defenders’ responses and escalating attack techniques accordingly (HackingTheUniverse). The problem with APTs is that they are—well, persistent. If the bad guys want something inside a given network badly enough, they will find a way to get in eventually.

Rule #2: Network defenders cannot change rule #1.

Despite the undeniable truth of Rule #1, network defenders have a hard time letting go of perimeter-defense strategies because they (the defenders) do not think about adversary goals beyond simply gaining network access. In reality, the bad guys want to do much more than just get into network systems; they want something, usually data that ranges from intellectual property for espionage purposes to personal data (credentials or personal identifiable information [PII]) for criminal financial gain. Better network defensive strategies would focus on how to keep adversaries from achieving their mission goals rather than just on how to keep them out. Such strategies could, for example, include hiding, disguising, or encrypting data so that attackers cannot find sensitive information. In addition, defenders should consider the possible benefits of mitigating the bad guys from inside a network instead of at the perimeter. In a perimeter defense, so little is learned about the adversary from the encounter. If an adversary is detected accessing a network, defenders need not let the attackers know they have been detected. By not signaling awareness of attackers, defenders stand a better chance of being able to see them the next time they attack (Rule #4 applies.).

Rule #3: They are already in.

While it is possible that the attackers have not gained network access, it makes good sense to assume that bad guys and their malware are already inside. Designing system-security controls and policies with the idea of limiting the damage an insider threat could achieve goes a long way toward mitigating the damage *all* threats can do, including those accessed remotely. A common approach to defending networks this way is the notion of 'least privilege'. Essentially, a user, program, or process is only allowed access to the minimum system resources required to perform legitimate functions. Also, designing for defense is not just about configuring network security-system devices and policies to look for internal threats. System users who assume their network is already compromised are suspicious of everything that looks out of place. The key element in all successful hacker tradecraft is the exploitation of trust, so the most effective network defenders do not trust anything that seems even a little bit peculiar.

Rule #4: Attacks will continue.

Defenders must avoid complacency. Just because today's attack was detected and turned away does not mean it is time to relax. Also, network defenders should not associate security compliance or good network hygiene with achieving good network security. To be sure, keeping up with all the latest virus definition updates, installing all the latest security patches, and reading all the security bulletins are prudent, but these measures will neither eliminate threats nor keep persistent attackers out. Constant vigilance is necessary, no matter how well protected a network might seem. (Rule #3 still applies.)

Rule #5: It is going to get worse.

It is safe to say that as technology is innovated and efficiency is increased through automation and artificial intelligence methods, so, too, will adversaries use the same innovations to increase the efficiency of their attacks. Today's attackers are often highly motivated and resourced; and as long as it continues to be profitable, attackers will go to great lengths to develop cutting-edge technology to break into networks.

Rule #6: Network defenders will contribute to worsening conditions.

In many cases, new cybersecurity capabilities actually accelerate the attack evolution and innovation. It is this rule that makes the best case for the need to understand co-evolution. In the bio-medicine domain, the medical community is revising its treatment strategies for viruses by using less aggressive measures for younger patients in order to avoid the eventual drug-resistant virus mutations. Cyber-defense strategists may want to consider whether there are analogs in the cybersecurity domain. (The Consequences of Cyberspace section below is relevant to such considerations.)

Rule #7: The only security achieved in the cybersecurity community is job security.

One might call this the ‘inconvenient truth’ about this domain, but it does not mean there has not been success in cybersecurity. There is also no clear definition of what makes a cybersecurity professional; there is no one-size-fits-all skill set. Despite the absence of a clear description of a cybersecurity professional, there is a pressing need for more of them.

So What Is This Co-evolution Thing?

The term ‘co-evolution’ is most often associated with the biological domain and typically refers to the natural adaptations species make, often influenced by natural selection, in order to survive in a given ecosystem. A prevalent type of co-evolution is that which is seen in predator-prey evolution where both sides evolve in terms of speed, stealth, camouflage, sense of smell, sight, and hearing as necessary to survive (for example, the polar bear is white to avoid being noticed when hunting, while the baby seal is also white to avoid being noticed by the polar bear).

So while co-evolution is primarily a biological concept, this phenomenon has been applied to other domains, including technological ones, by analogy. For example, computer software and hardware can be considered as two separate components, but they are tied intrinsically by co-evolution (D’Hondt *et al.* 2002). This idea is closely related to the concept of ‘joint optimization’ in socio-technical systems analysis and design. This kind of co-evolution can be characterized as mutualistic evolution; certainly, cyberspace activities have benefitted from this process. This article, however, explores the cybersecurity technological co-evolution that takes place between cyber attacker and defender, a process which clearly follows the predator-prey model more closely. Understanding this form of co-evolution enables defenders to position themselves strategically to get ahead of cyber threats.

A non-cyber example of co-evolution

Before exploring the cybersecurity domain, it might be helpful to show the attack-defend co-evolution phenomenon as seen from analogous examples in the transportation security domain. In just a few months following the events of 9-11, there was a failed attempt by a would-be terrorist to ignite explosive material concealed in his shoe—the infamous ‘shoe bomber’. The Transportation Security Administration (TSA) began requiring passengers’ shoes to be removed at airport security checkpoints so they could be X-rayed. In 2006, in an apparent response to checkpoint security measures, a terrorist plot to detonate liquid explosives was revealed, which resulted in another new security policy that banned liquids from carry-on baggage. This defensive measure was followed by another terrorist’s failed attempt to detonate explosives on a plane by hiding the explosive material in his undergarments (the ‘underwear bomber’) in 2009. Around the same time, the TSA had begun deploying full-body scanner devices at most major airports in the U.S. Perhaps in reaction to this security measure, terrorists attempted to hide explosive devices in printer cartridges being transported through express shipping services in 2010.

In hindsight, these examples illustrate the tendency of attackers to adapt in response to defensive measures and to maintain a consistent focus on their mission—in these examples, to get explosives on a plane where they could be detonated during flight. Understanding the attacker’s mission in a larger context is important because it may suggest constraints on an otherwise perceived infinite set of future attack adaptations. The notion of mission constraints will be explored in more detail later on in this article.

Evidence of Cyber Co-Evolution?

Likewise in the cyber domain, there are some illustrative examples of the attack-defend co-evolution phenomenon. For instance, in 2004, Microsoft released Service Pack 2 of its XP operating system that turned on its bundled firewall by default and included a new Data Execution Prevention (DEP) security feature (Microsoft 2004). The DEP feature provided protection against buffer overflow attacks (a popular infection vector for hackers), and some believe that the presence of this feature led hackers to move more toward file-format exploits against common desktop products, such as Adobe PDF and Microsoft Office documents. Similarly, after the Department of Defense (DoD) implemented Common Access Card (CAC) PKI authentication, considered a cybersecurity ‘game-changer’, many observed that malicious actors simultaneously increased their use of socially-engineered infection vectors. Another example of cyber co-evolution occurred when peer-to-peer (P2P)-based command and control (C2) botnets, such as ‘Storm Worm’, emerged shortly after the high-profile prosecution of some bot herders (FBI 2007) who used more centralized and attributable Internet-Relay-Chat- based (IRC) C2 mechanisms (TechShout Internet News 2007). Perhaps an even more direct correlation between cyberattack-defend occurred when the authors of the Conficker worm quickly adapted to a Microsoft-led cabal that attempted to pre-register and to lock out all of the worm's 250 pseudo-randomly-generated domains by rewriting the Conficker code to then generate fifty thousand domains for its update function (Keizer 2009).

Cybersecurity Community: Slow Evolution up to Now

For years, the network security community has been in a continuous struggle with malicious network attackers, constantly plugging holes in a very porous perimeter where defenders can only seem to see the holes after something has sneaked through or has leaked out. The developers of user applications have not been much help either, as they historically have put a priority on the users’ experience on the Internet rather than on their security. Although security features are becoming a priority to software vendors (of course, as security becomes more important to customers), new network security improvements still seem to lag far behind development of new attack methodologies (Jackson 2011).

Despite the best efforts of network defenders, no matter how much preparation went into the defense of the network, the conventional defender still waited, as if in a fortress, for the next breach of the virtual walls, not really knowing where an attack might come from and only being able to respond after the attack occurred. The advantage always seemed to be with the attacker, and the defender’s best case scenario was that the attack might be discovered before significant damage occurred. The point was (and still is) that the attacker or intruder always achieved some measure of success in every engagement.

Even now, a great deal of the emphasis in cybersecurity continues to be on threat detection where defenders seek to patch vulnerabilities and update sensors as quickly as possible after a potential threat has been discovered somewhere. While new advances in cloud technologies, heuristics, virtual sandboxing, and dynamic event processing have helped to decrease the shelf life of new attack vectors (for example, 0-day exploits), the net defenders are still operating very much in a react mode to an ever-increasing number of cyber threats. While getting better at knowing what can be known about threats, network defenders are still woefully ill-prepared for the threats that have not been discovered yet. Unfortunately, most discoveries take place long after the attackers have infiltrated victim networks and have achieved their mission goals. The best that a network defender can possibly achieve in this environment is to be protected against every form of attack short of those that employ 0-day exploits.

In recent years, with the advent of such organizations as the NSACSS Threat Operations Center (NTOC), U.S. Cyber Command (USCYBERCOM), FBI's National Cyber Investigative Joint Task Force (NCIJTF), and DHS National Cybersecurity and Communications Integration Center (NCCIC), and with the support of the intelligence community and commercial cybersecurity providers (such as McAfee, Mandiant, Symantec, Kaspersky, and others), network defenders now have much greater visibility into the human dimension of the cyber-threat environment. Cyber-threat analysts are beginning to delve deeper into the human behaviors behind the cyber-threat personas, even applying human science disciplines to the analysis of certain cyber-threat activities. As a result, the cybersecurity community is in a better position to anticipate some threat activities and to implement proactive defensive capabilities that go beyond the traditional, and mostly reactionary, perimeter-defense model.

Consequences of Cybersecurity (or Murphy's Law of Cybersecurity)

In the cybersecurity community, when defenders think about consequences, they usually think in terms of the consequences resulting from the attacker's actions and not the actions that they, the defenders, take. The defenders also tend to think of consequences in the short term and devote little energy to thinking about the longer-term effects of actions taken. One of the long-held goals of cybersecurity is to avoid strategic surprise, which typically comes from erroneous threat assessments and which results in the inability to anticipate a serious threat to an organization's vital interests. One source of strategic surprise actually comes from the unintended consequences that follow the implementation or adoption of certain cyber-defense measures along with the false sense of security perceived from those same measures. Just as there are rules for cybersecurity (described above), so there also seem to be innate rules (or laws) that reflect the cybersecurity community's current thinking about the consequences of cybersecurity actions. Below are seven rules about cybersecurity consequences that are intended to motivate defenders to think more strategically about cyber defense in a co-evolutionary context.

Rule #1: Every cybersecurity action has consequences; so does inaction.

As is true for other applications, short-term consequences are generally easy to see, but long-term consequences are much harder to predict. The rule seems to imply that failure to take an action (inaction) can result in bad consequences, but that is not necessarily true in all cases. At any rate, cybersecurity is never a long-term consequence of the action(s).

Rule #2: For each cybersecurity action considered, there are good consequences and bad consequences.

In the cybersecurity world, there are almost always tradeoffs, and as long as cybersecurity cost-benefit remains difficult to quantify, accurate consequence analysis (or risk assessment) will be very elusive. Most tradeoffs are seen in the context of security measures that can be implemented without creating a bad consequence for business/mission operations or user efficiencies. Another part of the problem is that most network defenders believe that all security measures result in good consequences.

Rule #3: Defenders usually only think about the good consequences.

In cybersecurity, many actions taken to protect or to defend a network are in response to an immediate problem and are initiated in order to generate a good consequence (that is, to remove the immediate problem). The fallacy with this kind of thinking is analogous to trying to douse an electrical fire with water: the reaction may be instinctive, but it is also ineffective, at best.

Rule #4: If the justification for supporting a decision to implement a particular cybersecurity action begins with the words, ‘The worst that could happen is …’, the network defenders probably have not really imagined *the worst* that could happen.

This line of thinking often occurs when defenders are uncertain about the resulting effectiveness of any given defensive action(s). When unsure whether something will work or not, defenders will often try to calculate worst-case scenarios; however, they may also forget that uncertainty works both ways in estimating best and worst outcomes.

Rule #5: When fighting a losing battle, the potential bad consequences of cybersecurity actions are usually forgotten.

Desperation will usually increase defenders’ risk appetite (rules 3 and 4 apply), especially if the defenders believe their own security (job security) might be at risk. If everything around them is on fire and water is all they have, they will use it, no matter the risks.

Rule #6: There is no credit for having healthy organs in a cadaver.

This rule may be viewed as a twist on the cliché ‘you’re only as strong as your weakest link’, but it is more likely related to the idea that one ‘Ah, sh*t!’ wipes out a dozen cases of ‘Well done!’ The weakest link is not always within the control of the net defender, but this does not mean it can be ignored. Defenders must be wary of the level of trust afforded to network affiliates whose cybersecurity posture might be inferior to their own. Also, overreliance on high-end technology solutions can often give defenders a false sense of security that may lead to even greater damage from compromise.

Rule #7: Defenders do not understand that the important thing is not being able to predict the consequence of actions; the most important thing is knowing all the potential consequences.

Meteorologists often have difficulty predicting the path of a storm too many days out, so they will usually present a set of potential paths based on different models and then monitor for weather conditions or indicators that favor a particular modelled path. By giving advanced notice of possible paths, civil emergency-management personnel have more time to take appropriate precautions in advance of the storm. Likewise in cyber defense, understanding the potential consequences and determining a means to monitor those consequences will allow defenders to be more proactive and to be better able to deal with consequences by developing appropriate, more resilient defenses.

Leveraging Attack-Defend Co-Evolution

The challenge is two-fold. First, defenders must learn to take advantage of this attack-defend co-evolution phenomenon in order to be more predictive about how attackers will respond to cyber defenses. Second, and more importantly, network defenders must devise methods to influence the attackers' evolution in a direction that plays into a position of strength for cybersecurity capabilities. In essence, there must be a shift away from a Tic-Tac-Toe network defense mentality, where the objective is more about trying not to lose, to a Chess game model where the best players are the ones who think several moves ahead. But how can such a shift be effected? Perhaps the first step is to determine what all the analogous chess pieces of a cyberattack/defend engagement are. Instead of looking at cyberattacks according to their individual technical components, defenders must view the attacks more holistically as (attack) systems that contain an arsenal of tools and techniques used to support different facets of an adversary's mission.

By analogy, when viewed as a weapon system, a military tank is seen as more than simply a big artillery gun that moves; it is also regarded as a system with armor for protection, small arms for self-defense, GPS and steering components for navigation, fueled engine and caterpillar tracks for mobility propulsion, camouflaging for stealth, communications for command and control, and radar and turret for targeting. Viewing cyberattacks as attack systems allows defenders to see components used to support similar needs, such as self-defense, propagation, stealth, command and control, and even striking a target. Too often, malware that performs multiple attack functions is abstractly characterized by defenders as one component of an attacker's arsenal. It is important to understand that not all attack functions need to evolve as a result of a given defensive action; in fact, the only ones that must evolve are the ones necessary to overcome or to circumvent the defensive measure that is inhibiting the attackers from completing their mission.

Slow the Attacker's Evolution with a Moving-Target Defense

The majority of conventional network defense models involve the use of mostly static tools and configurations. The problem with these defensive models and their static nature is that they are easily learned by malicious actors and thus allow attackers to rapidly adapt their attack methods and tools. Even so-called defense-in-depth and dynamic-defense capabilities can be learned by an attacker if used in a consistent manner over time. New cyber-defense strategies are calling for defenders to 'out-maneuver' attackers, which implies that defenses need to be able to maneuver or move. In this context, the concept of Moving-Target Defense (MTD) potentially comes into play. The Federal Networking and Information Technology R&D (NITRD) working group defines Moving-Target research as technologies that will enable defenders to "create, analyze, evaluate, and deploy mechanisms and strategies that are diverse and that continually shift and change over time to increase complexity and cost for attackers, limit the exposure of vulnerabilities and opportunities for attack and increase system resiliency" (NITRD 2010).

In almost any form of conflict or battle, a moving target is usually harder to hit, and if that target can move in an unpredictable manner, the attackers might expose themselves. The early notions of MTD were to be able to randomly move network nodes (timing and configuration) in order to impair an attacker’s reconnaissance of a target network and associated attack preparations (Okhravi *et al*. 2013). This initial approach actually posed new challenges for efficient (and cost-effective) management of network resources and potentially introduced new vulnerabilities.
Recently, the notion of MTD has been greatly expanded to include moving actual defensive devices, shifting defensive strategies, virtualizing, and creating various forms of deception. Now any kind of actions taken that can make the defender’s network less predictable would qualify as moving-target defense.

When paired with attack-defend co-evolution analysis, it might be possible to create an environment in which attackers are more predictable than the network defenses. Making network defenses less predictable makes it more difficult for attackers to learn and to adapt their attacks, thereby slowing the attack evolution. In this environment, the potential for the defender to influence an attacker’s movements in an advantageous direction should improve. Theoretically, MTD could be combined with aggressive cyber countermeasures, which, in the course of side-stepping an attack, puts the attacker in a vulnerable and exposed position that costs the attacker valuable resources. In other words, this model could increase the cost of business for a cyber adversary.

The following Spam behavior model provides an instructive example (Colbaugh & Glass 2011). **Figure 1**, below, is a graph of one of the key features of both spam and legitimate email. The graph shows an obvious and steady converging pattern of spam to legitimate email over a two-and-a-half-year period. Spam filters are trained to distinguish spam from legitimate email, and this convergence illustrates that a defensive strategy that does not change over time allows an adversary to learn and to adapt attack methods until the defensive measures become only marginally effective. An MTD strategy might have disrupted this learning process and slowed down the adaptation cycle.

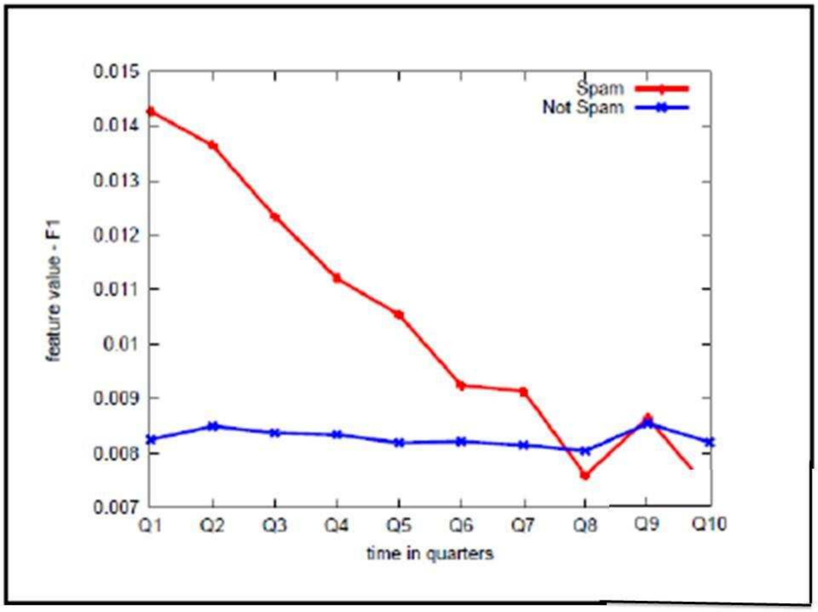

Figure 1: Spam vs. Non-Spam Features

It is important to note that the spammers' adaptions are constrained by their mission; that is, their spam must be understandable and must provoke a positive response from the spam recipient. The bottom line is that MTD probably offers the most opportunities to reduce the attackers' inherent advantages in cyber conflict and, ideally, even increases the chances that the defenders can get ahead of the attackers' operations' cycles.

Risk Assessment: Knowing What Attackers Want Is Key

Historically, network defense has been the job of system administrators schooled in the use of computer and network-security technologies. One of the big problems with this approach is that these people have limited understanding of the true vital assets of their organization and essentially work to defend the entire enterprise's IT system equally, which makes for a very large attack surface. It is a natural behavior for technical staff to think in terms of protecting the IT assets; but, in reality, the true goal is to protect the vital business processes, functions, services, and data that operate/reside in those IT assets—not necessarily the IT assets themselves.

To perform effective risk assessment, an organization needs to understand the threats of, vulnerabilities to, and consequences of compromise as they relate to the organization's particular business or mission. The risk analysis should begin with identifying an organization's vital assets, including such things as intellectual property, goods and services, customer/client data, and brand reputation. Once these assets are identified and quantified, then both vulnerabilities and threats can be more easily prioritized in the context of the consequences associated with the compromise of those assets.

Risk assessment in hand, defenders should actually now assume compromise (see Cybersecurity Rules 1, 2, and 3, above) and seek ways to mitigate damage from within the network. Knowing what a potential attacker is after, in other words, understanding the adversary's mission helps the defender understand the constraints the attacker might be under. This kind of knowledge should also help guide sensor placement and inform an MTD strategy that keeps an attacker from accessing the organization’s most important resources and assets. Ideally, the defender pushes the attackers away from sensitive areas of the organization’s network to areas of less concern, thereby increasing the attackers’ costs and lowering their return on investment.

Looking into the Past to Understand the Future

In order for the cybersecurity community to achieve a more holistic, mission-functional (chess pieces) system view of cyberattacks, it should look back at past events to see how defensive measures have influenced the evolution of attacks. Getting a sense of past adversary responses to cyber defenses could provide clues to the sensitivity of attack evolution to defensive actions, and could provide insights into how adversaries might respond to future cyber defenses. Borrowing from a similar study of how terrorists have responded to defensive technologies (Jackson 2007), cyber attackers could similarly respond to defensive measures in the following ways:

Altering operational practices

By changing the ways it carries out its activities or designs its operations, a cyber threat may blunt or eliminate the value of a defensive technology. Such changes frequently include efforts to hide from or otherwise to undermine the effect of the technologies

Making technological changes or substitutions

By modifying its own technologies (that is, exploits, encryption, malware, infrastructure), by acquiring new ones, or by substituting new technologies for those currently in use, a cyber-threat actor may gain the capacity to limit the impact of a technology on its activities.

Avoiding the defensive technology

Rather than modifying how cyber-threat actors blunt the value of a defensive technology, they might simply move their operations to an entirely different area to avoid the defensive technology. Such displacement changes the distribution of cyberattacks; and, although this may constitute successful protection in the area where the defensive technology is deployed, the ability to shift operations elsewhere limits the influence the technology can have on the overall threat level.

Attacking the defensive technology

If appropriate avenues are available, an attacker may seek to destroy or to damage a defensive technology to remove it as a threat or to turn the defensive technology into a greater vulnerability. For example, the implementation of Private Key Infrastructure (PKI) and digital certificates/digital signing were previously heralded as examples of game-changing technologies. Looking back at the adversaries’ evolution since the introduction of these technologies reveals that their responses have fit into all four of the aforementioned categories, although not all at the same time or in the same order. Adversaries are now at the point of attacking PKI and digital signing (attacking the defensive technology), and network defenders’ reliance on these technologies now potentially put them at greater risk than before these technologies were implemented. The moral of this story is that before defenders implement a ‘game-changing’ defensive technology, they must think about how the game is going to change (beyond the short-term effects).

Operational Targeting Cycle

Besides looking at attack-defend co-evolution from an attack-methodology perspective, another way to potentially observe co-evolution is by studying the attacker's targeting patterns. The defense actions and policies at a large enterprise or even at the sector level could influence the targeting patterns of certain classes of attackers. For instance, proactive cybersecurity policies and advances in fraud detection within the financial sector have likely affected targeting choices. In order to identify and to understand targeting patterns, defenders need to take a historical look at events, ideally those events where attacker groups can be distinguished. The defender should be looking for temporal patterns of activity in the context of socio-political and -economic events and in relationship to the business, mission, or function of the targeted network. Correlations may emerge between attackers, a category of events, and targets (victims) that might allow defenders to anticipate malicious activity as similar types of events occur. This model would operate much as the Center for Communicable Disease (CDC) quickly identifies likely disease strains during outbreaks based on incubation patterns observed in the past.

Through empirical evidence, network defenders can begin to anticipate when attacks are more likely to occur and will then have some insight into how attackers have previously responded to mitigation measures. This knowledge should allow network-security planners to gain the defensive 'high ground' and to implement mitigation actions more proactively, thus making their networks more resilient to attack.

By way of example, within the Department of Defense, there are component commands with specific regional and/or functional mission responsibilities that are potentially of high interest to foreign adversaries. These same commands perform regular, recurring training activities and exercises as well as respond to specific events or crises in their respective Areas of Responsibility (AORs), which are of interest to the nation's adversaries. It should also be no surprise that the networks of these units would be targets of foreign cyberattacks. If DoD network defenders can connect patterns of certain malicious activity with specific U.S. command operations and exercises, then adversary cyber activities become more predictable.

Using Models of Cyberattack-Defend Co-evolution for Cybersecurity Planning

Moving cyber-defense strategy planners from point-defense solutions to a defensive campaign mentality requires cyberattack system models that can provide the strategy planners with predictive analysis tools for more comprehensive cyber threat mitigation courses of action. Fortunately, there has been some recent foundational research modeling efforts that could support this cybersecurity planning approach. One such effort was a DoD-sponsored project entitled 'Cyber Adversary Dynamics', which had the specific goal of developing and demonstrating capabilities for modeling and exploiting co-evolution cyber behavior (Cybenko 2013). Among the key findings of the project were approaches to anticipating adversarial covert channel manipulations and a variety of approaches to defining the cyber 'high ground'.

Another promising DoD-sponsored research study from MIT Lincoln Labs developed a model of adaptive attacker strategy evolution and used it to investigate the strategies an attacker develops to overcome Moving-Target Defense strategies (Winterrose *et al.* 2014). Both this study and the previously mentioned DoD project used game-theoretic and adversarial reasoning approaches for model development. Finally, another related research effort from Mitre Corporation explored the notion of cyberspace deception and counter-deception as an emerging discipline (Stech *et al.* 2011). This work suggests that there is a need for more research on counter-deception in cyberspace to enhance security of computers and networks.

Managing the threat actor's experience

If effective predictive attack-defend co-evolution models can be produced, then it might be possible to use the insights gained from these models to conceive strategies for directing the attack evolution down paths that favor the cyber defenders, a process some call 'herding'. As previously described in the bio-medical example, the medical community is revising its treatment strategies for viruses by using less aggressive measures for younger patients in order to avoid the eventual virus mutations that are more resistant to prescription drugs. So, too, cyber defenders may not want to employ advanced, highly optimized security capabilities against all threats and in defense of all network assets. Instead, defenders may reserve these tools for the protection of their most highly valued assets. This approach may seem counterintuitive, since it requires defenders to accept some level of exposure of their networks to compromise. Perhaps a more effective strategy would be to employ a Moving-Target Defense strategy, as previously mentioned, that uses a mixed set of defensive measures to make it harder for attackers to learn the defense and to adapt their attacks.

Closing Thoughts

This article encourages the cybersecurity community to take a very strategic view of the cyber-threat environment and to look beyond immediate threat-response activities. Network defenders need to consider second- and third-order effects of their actions and proactively prepare for the next evolution of attack. To support this long view of the cyber threat, the cybersecurity community must develop models, tools, and technologies to help defense planners gain insight into evolutionary attack patterns and to avoid unintended consequences. In order to truly take advantage of these insights, the defenders must have some element of maneuverability within cyberspace; therefore, incorporating some form of a Moving-Target Defense strategy is strongly encouraged. Defenders should also consider some forms of deception or other techniques to create uncertainty for attackers. In essence, the cybersecurity community must find or develop ways to slow the attacker's evolution down since speeding up the defenders' evolution by very much is unlikely.

This article also makes some key assumptions about cyber attackers: first, they adapt to defenses; second, they are constrained by their mission goals; and third, this is not a zero-sum game. If these assumptions are valid, then cyber defenders should be able to take advantage of the co-evolutionary nature of cyberattacks and defend them to get ahead of attackers. Gaining this advantage will require the cybersecurity community to enlarge the depth of its knowledge of attackers, their mission goals, and the constraints associated with that mission. Fully leveraging this knowledge may require a fundamental shift in the way analysts view, describe, and document cyberattacks, as well as how they discern adversary capability and intent. It may also require fundamental shifts in the way network defenses are configured. Ultimately, in order to get ahead and to stay ahead of cyber threats, defenders need to decrease the predictability of their networks, increase the predictability of the attackers, influence the attackers' activities in a manner advantageous to the defenders, and avoid defensive strategies that produce unintended (negative) consequences.

Finally, before cyber and cybersecurity became part of the lexicon, 'Information Operations' was the concept most often associated with this problem space. This terminology was used almost exclusively in the military to describe the integrated employment of electronic warfare, computer network operations, psychological operations, deception, and operations security (DoA 2014). As public, private, government, and military global network infrastructures became interconnected in order to become this thing now called cyberspace, the term cybersecurity was coined, in part, as a more publicly palatable, non-military-centric concept. The downside of this universally recognized term is that it has perhaps unintentionally over-emphasized the technical, infrastructure side of the problem. In this context, another implicit assertion of this article is that it is essential for the community to return to thinking about cybersecurity more in terms of information assets and associated services since this is ultimately the commodity that most needs protecting.

References

Colbaugh, R & Glass, K 2011, 'Proactive defense for evolving cyber threats', *Proceedings of the 2011 IEEE ISI Conference*, viewed 17 March 2011, <http://ieeexplorer.ieee.org/stamp/stamp.jsp?tp=&arnumber=5984062>.

Cybenko, G 2013, *Cyber adversary dynamics*, viewed 30 December 2014, <http://www.dtic.mil/get-tr-doc/pdf?AD=ADA579667>.

Department of the Army 2014, Joint Publication 3-13, *Information Operations*, Chairman Joint Chief of Staff Publications, viewed 19 December 2014, <http://www.dtic.mil/doctrine/new_pubs/jp3_13.pdf>.

D'Hondt, T, De Volder, K, Kim Mens, K, & Wuyts, R 2002, *Co-evolution of object-oriented software design and implementation*, vol. 648, part 2, pp. 207–24, DOI: 10.1007/978-1-4615-0883-0_7.

FBI News Stories 2007, *Operation: Bot Roast 'Bot-herders' Charged as Part of Initiative*, viewed 4 June 2014, <http://www.fbi.gov/news/stories/2007/june/botnet_061307>.

HackingTheUniverse, *Advanced Persistent Threat*, viewed 13 January 2015, <http://www.hackingtheuniverse.com/infosec/isnews/advanced-persistent-threat>.

Jackson, B, Chalk, P, Cragin, R, Newsome, B, Parachini, J, Rosenau, W, Simpson, E, Sisson, M & Temple, D 2007, *Breaching the fortress wall: understanding terrorist efforts to overcome defensive technologies,* viewed 20 November 2013, <http://www.rand.org/pubs/monographs/MG481.html>.

Jackson, W 2011, ‘Clarke: Outdated cyber defense leaves US open to attack’, *Government Computer News*, viewed August 2012, <http://gcn.com/articles/2011/09/19/richard-clearke-us-outdated-cyber-defense.aspx>.

Keizer, G 2009, ‘Hackers update Conficker worm, evade countermeasures’, *ComputerWorld Security News*, viewed March 2009, <http://www.computerworld.com/s/article/9129239/Hackers_update_Conficker_worm_evade_countermeasures>.

Microsoft 2004, *Here’s why SP2 is such an important update for Windows XP*, viewed December 2014, <http://technet.microsoft.com/en-us/library/bb457009.aspx>.

Networking and Information Technology Research and Development (NITRD) 2010, ‘Moving target’, *NITRD CSIA IWG: cybersecurity game-change research & development recommendations*, 13 May, pp. 3-6, viewed 4 June 2014, <https://www.nitrd.gov/CSThemes/CSIA_IWG_Cybersecurity_Game-Change_%20RD_Recommendations_20100513.pdf >.

Okhravi, H, Rabe, M, Mayberry, T, Leonard, W, Hobson, T, Bigelow, D & Streilein W 2013, *Survey of cyber moving targets*, MIT LL Technical Report, Department of Defense, under Air Force Contract, viewed 31 December 2014, <http://www.dtic.mil/get-tr-doc/pdf?AD=ADA591804>.

Stech, F, Heckman, K, Hilliard, P & Ballo, J 2011 ‘Scientometrics of deception, counter-deception, and deception detection in cyber-space’, *PsychNology Journal*, vol. 9, no. 2, pp. 79-122, viewed 30 December 2014, <http://www.academia.edu/1328446/Scientometrics_of_Deception_Counter-deception_and_Deception_Detection_in_Cyber-space>.

TechShout Internet News 2007 *McAfee Avert Labs' top 10 security threats for 2008 announced*, viewed 14 May 2012, <http://www.techshout.com/internet/2007/17/mcafee-top-10-security-threats-for-2008-announced/>.

Winterrose, M, Carter, K, Wagner, N & Streilein W 2014, *Adaptive attacker strategy development against moving target cyber defenses*, Department of Defense, viewed 30 December 2014, <http://arxiv.org/pdf/1408.0023.pdf>.

Changing the Future of Cyber-Situational Awareness

N Newmeyer

Information Assurance Operations
National Security Agency, Fort Meade, Maryland, United States
JIWfeedback@nsa.gov

Abstract: *The proliferation of Internet of Things (IoT) devices will change the face of cyber situational awareness from one focused on centralizing and homogenizing data feeds to one struggling to identify triggers from inordinate amounts of data. IoT devices, anticipated to grow to 20-40 billion by 2020, will both increase the potential visibility and granularity of cyber situational awareness and will significantly complicate the effort. The sheer increase in communications will raise the noise floor and will force more advanced analytics and data parsing to identify appropriate triggers. In addition to the influx of data and traffic, IoT devices also have the potential to introduce server security concerns to any network.*

Keywords: *Internet of Things, Cyber, Situational Awareness, Sensor*

Introduction

Call it what you will: Internet of Things, Cyber Physical Systems, Pervasive Systems. They are all labels for the continued movement toward adding communication capability to items that historically would have had none. There is currently no indication that this trend will stop in the foreseeable future; it provides additional revenue potential for many commercial industries, and the public is perceived as being enamored by the capability to control everything through smartphones or tablets. Many of these devices, however, have little to no direct interface with the consumer; these intermediary sensors are designed to transmit information to control or network status systems. Data from both the consumer devices and the intermediary sensors will force significant changes in current network monitoring and situational-awareness capabilities.

Cyber-situational awareness has been described in a number of different ways; one of the cleanest explanations (Barford *et al*. 2010) separates it into three phases: situation recognition, situation comprehension, and situation projection. In this model, situation recognition encompasses awareness of the current situation, awareness "of the quality of the collected … information items," and "plausible futures of the current situation" (Barford *et al*. 2010). Situation comprehension focuses on awareness of the impact, awareness of actor behavior, and awareness of the causes of the situation; situation projection is entirely centered on the evolution of the situation (Barford *et al*. 2010). Leveraging this framework will enable a thorough discussion of the impact of increasing Internet of Things (IoT) devices on the future of cyber-situational awareness capabilities.

Current IoT Capabilities

The IoT is a concept that has been around for years, but it is being re-scoped to address the changing commercial market. While the concept has existed since 1991, the term has been in use since 1999 (Mattern & Floerkemeier 2010); the original concept focused on how Internet connected devices would change daily life through eliminating time-consuming functionalities, such as inventory control (Associati 2011). The current understanding focuses around interconnectivity of embedded devices that goes beyond Machine-to-Machine (M2M) communications (Holler *et al*. 2014) and is anticipated to drive automation in all associated fields, as well as to create new commercial opportunities. The variety of capabilities that fall within the current IoT spectrum ranges from smart meters to tire pressure sensors to heart monitors to vending machines (Wigmore & Rouse 2014).

Beecham Research has designed a sector map (**Figure 1)** dividing the world of IoT into what it describes as service sectors and providing examples of IoT devices that fall into each category. The diagram below illustrates perfectly the breadth of IoT and the potential consumer impacts; each of the categories of devices listed around the outside of the diagram represents an entire commercial market sector: surgical equipment, environmental monitors, HVAC, vehicles. Each of those market sectors has its own communications requirements, security concerns, and functional needs. To date, the individual industries responsible for each sector have had primacy in defining the requirements for their components.

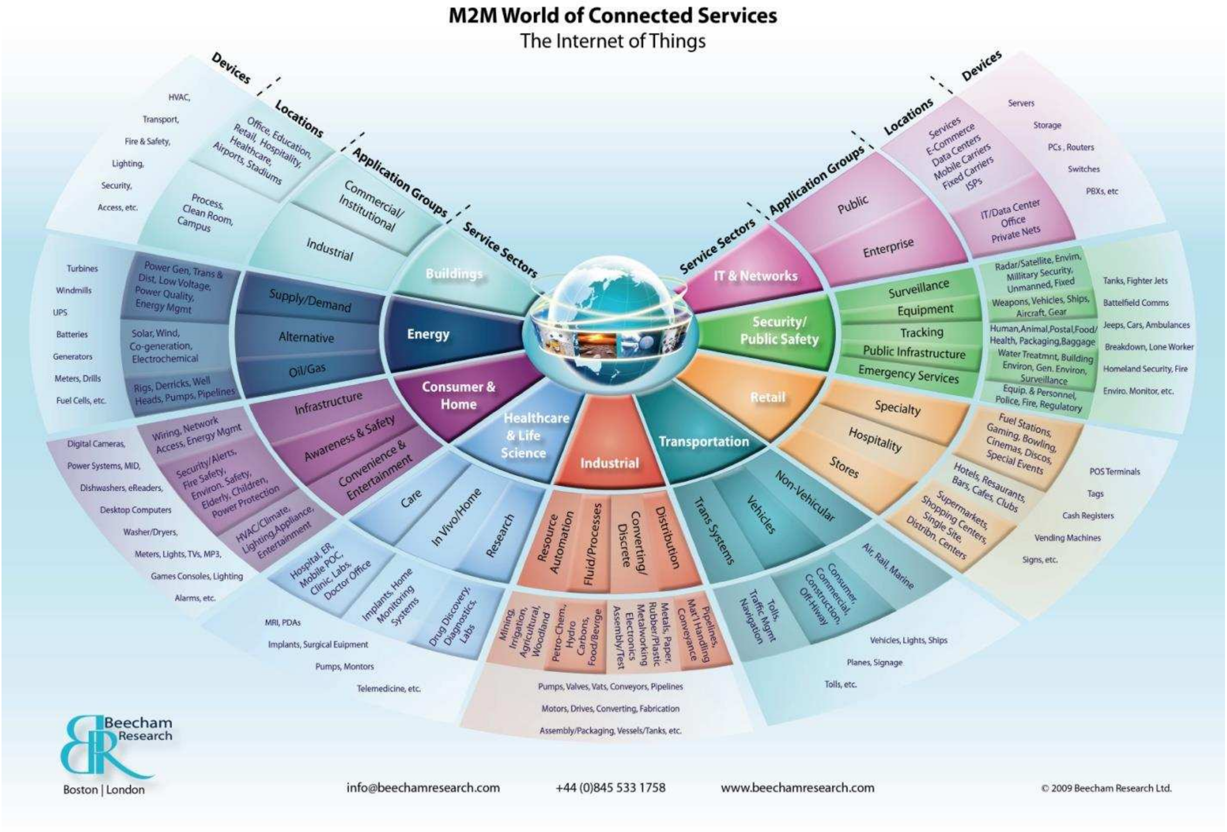
M2M World of Connected Services
The Internet of Things
Devices
Locations
Application Groups
Service Sectors
HVAC,
Transport,
Fire & Safety,
Lighting,
Security,
Access, etc.
Office, Education, Retail, Hospitality, Healthcare, Airports, Stadiums
Process, Clean Room, Campus
Commercial/ Institutional
Industrial
Buildings
Turbines
Windmills
UPS
Batteries
Generators
Meters, Drills
Fuel Cells, etc.
Power Gen, Trans & Dist, Low Voltage, Power Quality, Energy Mgmt
Solar, Wind, Co-generation, Electrochemical
Rigs, Derricks, Well Heads, Pumps, Pipelines
Supply/Demand
Alternative
Oil/Gas
Energy
Digital Cameras,
Power Systems, MID,
Dishwashers, eReaders,
Desktop Computers
Washer/Dryers,
Meters, Lights, TVs, MP3,
Games Consoles, Lighting
Alarms, etc.
Wiring, Network Access, Energy Mgmt
Security/Alerts, Fire Safety, Environ. Safety, Elderly, Children, Power Protection
HVAC/Climate, Lighting,Appliance, Entertainment
Infrastructure
Awareness & Safety
Convenience & Entertainment
Consumer & Home
Healthcare & Life Science
Care
In Vivo/Home
Research
Hospital, ER, Mobile POC, Clinic, Labs, Doctor Office
Implants, Home Monitoring Systems
Drug Discovery, Diagnostics, Labs
MRI, PDAs
Implants, Surgical Euipment
Pumps, Montors
Telemedicine, etc.
Industrial
Resource Automation
Fluid/Processes
Converting/ Discrete
Distribution
Mining, Irrigation, Agricultural, Woodland
Petro-Chem., Hydro Carbons, Food/Bevrge
Assembly/Test Electronics Metalworking Rubber/Plastic Metals, Paper,
Pipelines, Mat'l Handling Conveyance
Pumps, Valves, Vats, Conveyors, Pipelines
Motors, Drives, Converting, Fabrication
Assembly/Packaging, Vessels/Tanks, etc.
Transportation
Trans Systems
Vehicles
Non-Vehicular
Tolls, Traffic Mgmt Navigation
Consumer, Commercial, Construction, Off-Hiway
Air, Rail, Marine
Vehicles, Lights, Ships
Planes, Signage
Tolls, etc.
Retail
Specialty
Hospitality
Stores
Fuel Stations, Gaming, Bowling, Cinemas, Discos, Special Events
Hotels, Resaurants, Bars, Cafes, Clubs
Supermarkets, Shopping Centers, Single Site, Distribn. Centers
POS Terminals
Tags
Cash Registers
Vending Machines
Signs, etc.
Security/ Public Safety
Surveillance
Equipment
Tracking
Public Infrastructure
Emergency Services
Radar/Satellite, Envirn, Millitary Security, Unmanned, Fixed
Weapons, Vehicles, Ships, Aircraft, Gear
Human,Animal,Postal,Food/ Health, Packaging,Baggage
Water Treatmnt, Building Environ, Gen. Environ, Surveillance
Equip. & Personnel, Police, Fire, Regulatory
Tanks, Fighter Jets
Battelfield Comms
Jeeps, Cars, Ambulances
Breakdown, Lone Worker
Homeland Security, Fire
Enviro. Monitor, etc.
IT & Networks
Public
Enterprise
Services E-Commerce Data Centers Mobile Carriers Fixed Carriers ISPs
IT/Data Center Office Private Nets
Servers
Storage
PCs , Routers
Switches
PBXs, etc
Beecham Research
Boston | London
info@beechamresearch.com
+44 (0)845 533 1758
www.beechamresearch.com
© 2009 Beecham Research Ltd.

Figure 2: Beecham Research IoT Diagram (Beecham Research 2011)

As the IoT market has continued to grow, organizations have recognized the need for common standards and for a common framework to ensure both interoperability and data security. In July 2014, IEEE kicked off an effort to define these standards with industry collaboration and create a standard architecture to bring uniformity to the currently disjointed market (Lawson 2014). The 23 vendors participating in the group do not intend to replace any of the existing IoT groups, but they do intend to provide order and cooperation between the many standards bodies involved in these discussions. As of today, however, there is no unifying structure or security standard governing this wide range of capabilities; devices are not necessarily interoperable; personal data is protected differently and, potentially, in a manner that is not publicized. This situation may not be concerning to the public from the perspective of an Internet-connected washing machine, but a point-of-sale smartphone credit-card reader that has access to personal information is a different story. Potentially even more concerning to the public would be any security issues associated with an Internet-connected smoke alarm or carbon monoxide monitor, along with any privacy concerns associated with health information traversing these broad networks.

Predicted IoT Proliferation

Security concerns put aside for the time being, with the current scope of capabilities that are considered to fall within the IoT spectrum, one of the next logical questions involves how the commercial market is predicted to grow. Current market reports indicate there are 1.9 billion IoT devices in the hands of consumers today, and that number is predicted to grow to over 9 billion in four years (Adler 2014); Cisco research indicates IoT traffic will have an annual growth rate of 84% through 2018 (Cisco 2014). International Data Corporation (IDC) expands that prediction to 2020, stating that there will be 30.1 billion “connected autonomous things” (*Business Wire* 2013). Even if these predictions severely overestimate the proliferation of IoT devices, the actual increased numbers will change the current Internet environment.

IoT Security

In part due to the incredibly high predictions of market growth, concerns about the security of IoT devices are becoming more widespread. Beecham Research has published statements on existing IoT vulnerabilities and security flaws, to include the Stuxnet attack on Industrial IoT, and attacks on consumer-connected lighting products (Beecham Research 2014). Forbes has focused on the more sensationalistic approach to expanding consumer awareness, by summarizing potential attack surfaces within consumer IoT devices that could be vulnerable to hackers (Steinberg 2014). Awareness of security issues across consumer, business, and government user bases is vital to the success of IoT devices; the publicized logical connection between the possibility of cyberattacks and IoT will enable future discussions about embedded security mechanisms to proceed. Many academic and business institutions are calling for security mechanisms to be built into these devices from the ground up and not added as afterthoughts (Clearfield 2013a); this vital posture shift will prepare the IoT market for the anticipated growth and future spread of devices.

U.S. government organizations have considered the proliferation of IoT devices a serious security concern since 2008 (National Intelligence Council 2008), but have not yet adjusted policy to account for the integration of these devices in U.S. defense networks (Committee on National Security Systems 2014). The FTC has tried to levy sanctions on specific IoT vendors to ensure advertised security parameters are met, but the FTC does not regulate privacy or security and cannot extend its reach into those areas (Clearfield 2013b).

IoT Impact on Cyber-Situational Awareness

Current cyber-situational awareness limitations

Current cyber-situational awareness capabilities are limited in the visibility they provide into network traffic and incidents. Today’s techniques include a combination of vulnerability analysis, intrusion detection, and forensics, among others, to provide low-level situational-awareness information (Barford *et al*. 2010). Anything else is currently completely dependent on a human performing manual analysis; indeed, “[t]here is still a big gap between human analysts’ mental model and the capability of existing cyber situation-awareness tools” (Barford *et al*. 2010).

On top of that limitation, there are already complex analytical requirements based on the amount of data that exists today. Some systems have real-time data-processing requirements, which heavily influence the analytic frameworks developed; based on multiple sensors’ reporting data, there is regularly significant redundancy that existing analytics are not always equipped to handle. This redundancy introduces the potential for synchronization issues as well as questions about the veracity of the analysis (Barford *et al*. 2010).

The Beecham Research IoT sectors (**Figure 1**, above) can be viewed as those containing devices that interact directly with consumers and those containing devices that interact primarily with other devices. The Machine-to-Machine (M2M) sensors will have the most impact on cyber-situational awareness. These communications have been divided into four primary components (Chen 2012): data collection, information relay, data analysis, and services taking action based on analysis results.

Assuming a conservative assessment of the future spread of IoT devices, if, by 2020, there are 25 billion devices connected to the Internet, every one of the four components listed above will be impacted. If IoT devices are increasing at anywhere close to 84% per year, today's data-collection methodologies will need to dramatically change. Enterprise storage capacity will have to increase to account for the increased data, as well as throughput in every enterprise infrastructure node that handles IoT traffic. Information relay will also be affected: not only will the network capacity of infrastructure nodes have to be significantly increased, but the processing capability will also have to be able to perform the same analytic functionality at the same speed on exponentially more data. While deconfliction and redundancy are issues today, they will be even greater issues as IoT traffic increases. Analytics will need to perform at the same speed on repeatedly increasing data sets while ensuring the data is trustworthy. All of this processing will need to occur before any service designed to take action can perform. To handle these changes, most enterprises will likely be forced to upgrade their entire network-monitoring and event-storage infrastructure.

While much of this may sound negative, increased proliferation of IoT devices will bring unprecedented granularity and breadth of understanding to cyber-situational awareness; it will just require an initial investment from individual enterprises to support the increased data flow and analytic capabilities first. Incident-response support today relies on collection of data from host-based sensors and network-based sensors; that data often has to be correlated manually by analysts to determine the progression of a situation or speculate on situational futures. These predictions are limited based on the information the analyst has access to: if the host-based agent or network-based agent does not have complete access to the data relating to the potential compromise, the analyst is working in the dark.

Situation recognition

Of the two primary components of situation recognition, the first is awareness of the situation, including the ability to assess the quality of information and trustworthiness of the data that is being provided by the network. Gaps in data provided to the analytic platform will immediately result in an inaccurate assessment of the situation itself. In the most egregious of those scenarios, if an event were to prevent data from reaching the collection point, it is possible the enterprise would not be able to recognize the situation at all. However, if the network were enumerated with countless IoT devices, even as simple as ZigBee nodes within the infrastructure, it would make it virtually impossible to prevent all indications of an event from reaching the monitoring system. While analytics have to be developed to handle the additional data feeds from IoT devices within an enterprise network and the data correlated with the existing host-based and network device-based sensors, that additional data adds both depth to the situation-recognition capability and confidence that false negatives are at a minimum.

For an enterprise to truly have that additional confidence, its analytics would need to be strengthened to account for not only additional data from IoT devices, but also for an actor attempting to hide behavior in the now increased noise. The decentralized nature of IoT communications would require that analytic capabilities be adjusted to account for both the IoT communication types and the potential attempts of an actor to leverage the IoT devices themselves. These signatures would differ between consumer IoT devices and M2M IoT devices, requiring both to be accounted for in an enterprise's situational-awareness analytics.

The second primary component of situation recognition is assessing the plausible futures of the current situation. With incomplete or contradictory information, assessing future direction of a situation with any confidence is highly unlikely. Predictions can be made, and are made with regularity, at a high or broad level based on the information available today during an incident response. Additional sensor data resulting from numerous IoT devices being deployed throughout enterprise networks would provide increased clarification on situation specifics and would lead to increased confidence in futures predictions. In the event that an enterprise had not advanced its analytic capacity to ingest data from IoT devices within its network and to include that information in the situation assessment, the enterprise would not only be blind to an actor that leveraged those IoT devices, but their analytics would also not be mature enough to accurately predict futures of any situation. Once an actor advanced to the point of leveraging IoT devices anywhere in the capability suite, an enterprise would need to account for that tactical change within its network defenses.

Situation comprehension

Unlike situation recognition, which focuses on the initial identification of a situation and its plausible futures, situational comprehension is focused on awareness of impact, actor behavior, and the causes of the situation. This broader scope requires focus on multiple facets of the available data. Awareness of impact requires that the analytic can determine the current scope of the situation and can accurately assess what functionally has transpired. Actor behavior would be assessed slightly differently, by analyzing any indicators collected while an actor is active on the enterprise network as well as considering the focus and propagation of any initial compromise detected. To enumerate causes of the situation, the analyst must be able to determine how the actor initially gained access to the network as well as to identify any potential reasons for that behavior or for targets within the enterprise network. Any one of these tasks would be incredibly difficult in an optimal situation; with today's data, however, the analyst would be in a far-from-optimal situation. Most networks are not fully enumerated. Consequently, an analyst would have access to data from a minimal set of nodes; that data might not be archived anywhere or remote nodes might not have the same level of monitoring. Addressing those issues with minimal data on actor behavior would almost never result in a confident assessment.

With the addition of IoT devices within the enterprise, and enhanced situational awareness and analytic capability to support the data from the IoT devices, an analyst would have greater visibility into any network penetration and into any actor movements on the network. Without the integration of the IoT devices into the enterprise situational awareness and analytic capability, efforts to identify causes of a situation would always have a large blind spot. In the near future, it will be more and more difficult to prevent IoT devices from being integrated with enterprise networks.

Situation progression

Similar to the other categories, determining the evolution of the situation without a thorough set of data would only lead to a weak result. The more data the analyst has access to regarding the status of the computers and the network during the alleged event, the more accurate assessment of the situation progression that can be made. If the investigator does not have an accurate picture of the full enterprise network, including any integrated IoT devices, he or she cannot accurately predict how an actor's tactics are going to evolve or shift throughout a situation. Similarly, if an actor historically preferred presence on one type of machine, why would an analyst not review data for that type of machine?

This area is fraught with the same concerns as the others: higher throughput sensors, additional analytics, higher likelihood that the adversary could hide in the increased noise floor. But without looking forward to address those concerns, cyber-situational awareness will be in an even weaker state.

Visualization

Proliferation of data relating to incident-response actions and investigations of actor behavior will drive requirements for advanced visualization capabilities. Current cyber-situational awareness-visualization capabilities are tied to customized display algorithms developed for an enterprise- or instance-driven scripting performed at the time of need. Capabilities will have to be upgraded to account for the significant number of increased data feeds from IoT devices, as well as to design a way to ensure the display is usable and not covered in additional data.

IoT and Cloud

The significant increase in data volume that comes with IoT devices will lead to a discussion of big data or cloud storage if nothing else were needed. Many enterprises are already upgrading their storage infrastructure to a cloud configuration to support big data analytics. These structures are significantly beneficial to an IoT-integrated enterprise, and they significantly lessen the upgrades necessary for an enterprise to support. The large amounts of storage provided by a cloud architecture increases the likelihood that data from host-based and network-based sensors will be archived and available for later large data set analysis to be performed. This storage would also support deconfliction from different information sources; it would provide an opportunity for validating consistency of data sets and feeds, and would eliminate processing of duplicate data. Large data set analysis would also be possible, potentially searching for lesser occurrences that might not initially trigger concern. Each enterprise would have to assess whether it had any real-time analytic requirements, and how to shift any current processes to operate on IoT data stored in cloud storage.

If it were not feasible to update the architecture in a way that would enable the increased IoT traffic to flow back to the storage mechanism, an enterprise could assess mechanisms of adding processing and deconfliction to the remote sensor nodes—potentially minimizing the amount of data to transmit back to the storage mechanism, but requiring more processing capability on the remote nodes.

Looking Forward

Data from IoT devices embedded in enterprise networks will fill a vital role in improving and maturing cyber-situational awareness capabilities across industries. One key characteristic of that role is that the analysts and the enterprises can trust the data being returned from the IoT devices. Currently, that trust can vary depending on the specific security mechanisms included in the individual IoT devices by the manufacturer. There is no industry-wide standard set of security capabilities or trust mechanisms required for IoT devices.

Not only will the expanding IoT market impact commercial enterprises, but it will also impact all organizations with information technology infrastructure: from restaurant purchasing point-of-sale terminals to a government organization buying new servers. As the proliferation continues, it will become more difficult to prevent IoT devices from becoming integrated with an enterprise's networks; as the interconnectivity is an eventuality, the community needs to act now to ensure appropriate security and trust mechanisms are included in the devices.

Without standard security mechanisms, there is little to no trust in the validity of data, which impacts an enterprise's ability to recognize future implications of current situations and to accurately assess the impact of current incidents.

Security Needs

The state of current documentation on IoT component security is exceptionally tenuous. While the field has matured within the past few years, much discussion regarding security and security standards is still necessary. Most of the literature in this area has been published by commercial research groups and academics; recognized industry organizations such as IEEE are just starting to join the discussions, and the U.S. government has yet to publish policy or requirements that sufficiently address IoT devices.

Some authors have taken a sensationalistic approach; they are attempting to instill fear into the consumer and are expounding on the multitude of attack surfaces these devices bring into homes (Steinberg 2014). Others have grouped past public-network attacks into categories of IoT devices, to illustrate how complicated the environment has become and how vulnerable the world's infrastructure is (Beecham Research 2014).

Many academics are approaching the topic very similarly to IEEE (IEEE 2014; Lawson 2014), by recommending that security be built in to the design of all IoT components as a primary consideration, rather than being added as an afterthought (Clearfield 2013a). They uniformly argue that, if security is not a core component of these devices, the predicted market expansion (Cisco 2014; *Business Wire* 2013) has the potential to expose sensitive data in amounts never before seen.

The U.S. federal government has published a few reports referencing IoT devices, but nothing yet that addresses a way forward. In 2008, IoT was identified as a serious security concern (National Intelligence Council 2008); but not until 2013 were public and private companies engaged to explore the benefits and concerns of IoT through the SmartAmerica Challenge (Voyles 2014).

Following that, in 2014, the U.S. President's National Security Telecommunications Advisory Committee took an initial look at IoT in its Industrial Internet Scoping Report (National Security Telecommunications Advisory Committee 2014); the report identifies a need for a federal strategy for IoT and takes responsibility to create an initial strategy by November of 2014. While an initial high-level strategy may soon be available, that alone is not sufficient to quantify the impact of and requirements for IoT device interactions with U.S. defense networks.

Some authors are beginning to push harder for the inclusion of security mechanisms in IoT devices. Current discussions focus on identity validation, authentication improvements, and access control (Ndibanje *et al*. 2014). While no currently proposed security mechanism has significantly more market support than others, what matters today is that industry moves forward to support security mechanisms.

References

Adler, E 2014, *The 'Internet of Things' will be bigger than the smartphone, tablet, and pc markets combined*, viewed 30 September 2014, <http://www.businessinsider.com/growth-in-the-Internet-of-things-market-2-2014-2>.

Associati, C 2011, 'The evolution of Internet of Things', *Focus,* February 2011, viewed 30 September 2014, <http://www.slideshare.net/casaleggioassociati/the-evolution-of-internet-of-things>.

Barford, P, Dacier, M, Diettrich, TG, Fredrikson, M, Giffin, J, Jajodia, S, Jha, S, Li, J, Liu, P; Ning, P, Ou, X, Song, D, Strater, L, Swarup, V, Tadda, G, Wang, C & Yen, J 2009, 'Cyber SA: Situational Awareness for cyber defense', *Cyber Situational Awareness: Advances in Information Security,* 30 September, vol. 46, pp. 3-13, ISBN: 978-1-4419-0140-8, Springer, U.S.

Beecham Research 2011, *M2M world of connected services: Internet of Things*, Beecham Research, Boston, MA.

——2014, *IoT security study initial report: executive summary*, Beecham Research, Boston, MA.
Business Wire, 2013. 'The Internet of Things is poised to change everything, says IDC', *Business Wire*, Hathaway, Berkshire.
Chen, YK 2012, 'Challenges and opportunities of Internet of Things', *IEEE Design Automation Conference 2012, Sydney NSW,* pp. 383-88, ISBN: 978-1-4673-0770-3.

Cisco 2014, *Cisco visual networking index: forecast and methodology, 2013-2018,* Cisco, n.p.

Clearfield C 2013a, 'Rethinking security for the Internet of Things', *Harvard Business Review*, June 2013, viewed 28 September 2014, <http://blogs.hbr.org/2013/06/rethinking-security-for-the-in/>.

——2013b, 'Why the FTC can't regulate the Internet of Things', *Forbes*, viewed 29 September 2014, <http://www.forbes.com/sites/chrisclearfield/2013/09/18/why-the-ftc-cant-regulate-the-internet-of-things/>.

Committee on National Security Systems (CNSS) 2014, *CNSSP no. 17: Policy on wireless systems*, National Security Agency, Fort Meade, Maryland, United States.

Holler, J, Tsiatsis, V, Mulligan, C, Karnouskos, S & Boyle, D 2014, *From machine-to-machine to the Internet of Things: introduction to a new age of intelligence*, 1st ed., Academic Press, Elsevier, Oxford.

IEEE 2014, *IEEE Internet of Things*, viewed 30 September 2014, <http://iot.ieee.org/>.

Lawson, S, 2014, *IEEE standards group wants to bring order to the Internet of Things,* viewed 30 September 2014, <http://www.computerworld.com/article/2686714/networking-hardware/ieee-standards-group-wants-to-bring-order-to-Internet-of-things.html>.

Mattern, F & Floerkemeier, C 2010, *From the Internet of computers to the Internet of Things*, ETH, Zurich (Swiss Federal Institute of Technology Zurich), Institute for Pervasive Computing, Distributed Systems Group, viewed 29 September 2014 <http://www.vs.inf.ethz.ch/publ/papers/internet-of-things.pdf>.

National Intelligence Council 2008, *Disruptive civil technologies: six technologies with potential impacts on US interests out to 2025*, viewed 28 September 2014, <http://fas.org/irp/nic/disruptive.pdf>.

National Security Telecommunications Advisory Committee 2014, *The President's National Security Telecommunications Advisory Committee: Industrial Internet scoping report*, 19 February, NSTAC, Department of Homeland Security, n.p.

Ndibanje, B,. Lee, H-J & Lee, SG 2014, 'Security analysis and improvements of authentication and access control in the internet of things', *Sensors: Wireless Sensor Networks and the Internet of Things*, 13 August, pp. 14786-14805, DOI:10.3390/s140814786.

Steinberg, J 2014, 'These devices may be spying on you (even in your own home)', *Forbes*, viewed 3 October 2014, <http://www.forbes.com/sites/josephsteinberg/2014/01/27/these-devices-may-be-spying-on-you-even-in-your-own-home/>.

Voyles, R 2014, *SmartAmerica challenge: harnessing the power of the Internet of Things,* viewed 5 October 2014, <http://www.whitehouse.gov/blog/2014/06/10/smartamerica-challenge-harnessing-power-Internet-things>.

Wigmore I & Rouse, M 2014, *Definition: Internet of Things (IoT),* viewed 3 September 2014, <http://whatis.techtarget.com/definition/Internet-of-Things>.

The Need for Digital Identity in Cyberspace Operations

AR Friedman, LD Wagoner

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

Abstract: *Digital identity is an online or networked identity in cyberspace for an individual, organization, or entity to uniquely describe a person or a thing and contains information about the entity's relationships. A critical challenge in cybersecurity and cyberspace operations is knowing with whom or what one is defending. Currently, it can be difficult to accurately determine the identity of a person or entity in cyberspace. A unified and verified identification system for each entity or component of an IT system is needed. This paper will identify the challenges and opportunities that digital identity technologies introduce for cybersecurity and cyberspace operations.*

Keywords: *Digital Identity, Software ID Tags, SWID, NSTIC, IDESG*

Introduction

In the U.S. National Military Strategy for Military Operations, the information environment within cyberspace is described as "the aggregate of individuals, organizations, and systems that collect, process, disseminate, or act on information" (Joint Chiefs of Staff 2006). In order to conduct cyberspace operations and operate securely in cyberspace, it is essential that assets can be identified in real time. There is an old adage that states that one cannot manage what one cannot measure. However, those who are engaged in or with the current state of information technology (IT) (individuals and/or organizations) and even the systems themselves are at a more precarious point: they are not able to manage their IT systems because they do not even know what or who is on those systems. As a result, the need exists to appreciably ramp up the ability to identify both what and who are operating on any given IT system.

What Is a Digital Identity?

The United States Federal Chief Information Officer (CIO) Council has defined digital identity as "[t]he representation of identity in a digital environment" (Executive Office of the President of the U.S. 2011a), while the National Strategy on Trusted Identities in Cyberspace (NSTIC) defines it as "a set of attributes that represent a subject in an online transaction" (Executive Office of the President of the U.S. 2011b). Digital identity also has another common functional definition as "the digital representation of a set of claims made by one digital subject about itself or another digital subject" (Executive Office of the President of the U.S. 2009). All three of these definitions are from U.S. government publications. However, as will be clear, each does not adequately capture the general sense of the term 'digital identity'.

Technopedia offers a formal definition of digital identity as

ISSN 1445-3312 print/ISSN 1445-3347 online

> an online or networked identity adopted or claimed in cyberspace by an individual, organization or electronic device. These users may also project more than one digital identity through multiple communities. In terms of digital identity management, key areas of concern are security and privacy. ('Digital identity')

Although Technopedia's definition is broadly referenced in a variety of sources, the definition is restrictive in that it requires that the identity be adopted or claimed. This is not necessarily the case. Cognizance of this fact leads to a better definition:

> A digital identity is defined as a set of data that uniquely describes a person or a thing (sometimes referred to as a subject or entity) and contains information about the subject's relationships to other entities. (Windley 2005)

One of the main purposes of a digital identity is to enable the differentiation of an entity from a multitude of entities. For an IT system, unique means of identification are needed to differentiate entities. For example, a unique digital identity for a person could be an email address. For a corporate IT system that controls and monitors assignment and use of email addresses to its employees, an email address can serve as a reliable identifier of the person using the email account. However for many IT systems, there is minimal trust in the verification of the actual identity of the person who either establishes the email account or uses it, making an email address limited in usefulness as a digital identity.

U.S. and foreign government organizations, as well as the commercial sector worldwide, have recognized the need for more authoritative digital identities for people. For instance, criminals are falsifying tax returns using stolen social security and national identity numbers. As a result of these and other breaches of digital identity security, "[m]any European countries have been investing in national e-ID systems, as have countries in the Middle East and Asia" (Castro 2011). Despite these efforts, identity theft continues to grow worldwide, and new laws and technologies are needed to protect citizens and businesses. One response to these needs is taking shape in quite a few states in the U.S. that are implementing plans in their respective Departments of Motor Vehicles to enable identity services that will share authentication data between state agencies to improve service, to reduce fraud, and to cut costs.

Though there is an undeniably immediate need for developing strong digital identities for people, the need to identify and to track non-person entities is just as pressing. In some countries, this need is being addressed by the use of biometric and credentialing technologies that create digital identities that can identify and track dangerous cargo, which includes materials such as explosives and poisonous gas. Indeed, the use of digital identities to authoritatively identify entities in cyberspace has widespread application. As the Internet of Things (IoT) expands rapidly, the need for reliable digital identities for entities in cyberspace will become even more critical.

Digital Identity in an Information Technology System

All parts of an IT system residing in cyberspace will need digital identities. Though digital identity is often thought of in terms of being able to identify people in cyberspace, people are just one part of an IT system. In order to conduct cyberspace operations and sufficiently protect IT assets, policy makers and technologists need to understand that an IT system is a closely connected group of interconnected elements that all require digital identity. Specifically, an IT system is composed of five parts: people, procedures, software, hardware, and data (Silver, Marcus & Beath 1995). These components interact to form a system that is functional and responsive to the needs of its users. Procedures set the rules and guidelines that allow people to run software on a hardware platform to manipulate data to produce desired results. Moreover, each of these categories can have multiple subcategories. For example, software has two major subcategories: system software and application software. Data can be in many forms: such as raw unprocessed bits, document files, worksheet files, database files, and presentation files.

Of the five parts of an IT system, the number of procedures is relatively small and stable compared to the other four parts. The sheer magnitude of the numbers of the remaining four components in an IT system—along with their constant volatility—creates a significant challenge to identify and to manage those components. Creating a digital identity for each of the components' entities would enable them to be identified accurately, efficiently, and quickly even in a continuously fluctuating environment. Despite the fact that there are many diverse authentication systems and digital identifiers that attempt to address the problems of uniqueness and authenticity for an individual component of an IT system, much work is still needed to accurately identify and to track the individual components within an IT system.

Furthermore, being unable to identify, with precision, each entity within the remaining categories of people, software, hardware, and data in use is unacceptable and is contrary to effective military cyberspace operation and cybersecurity. Digital IDs for each of the entities within these four IT components can provide the leap ahead needed to identify what is on IT systems, how they are being used, and where they are being used. There are unique challenges and different identity techniques that are emerging to create digital identity within each category. While the challenges of each of these categories may be unique, they all share one need in common: to be able to identify individual entities in each respective component category.

As part of the maturation of IT systems, there are emerging initiatives to address the needs of items in each category and the needs they share in common. The following sections seek to identify these emerging initiatives.

People

As previously discussed, one form of digital identity for people is an email account. Though it is not currently an authoritative digital identifier in many instances, the email account is often used to form a 'hub and spoke' architecture for other accounts. The email account is linked to other accounts such that if the password is forgotten for any of the 'spoke' accounts, an email can be sent to the 'hub' account with instructions for resetting the password. Given the ease and lack of security in obtaining most email accounts, this protocol only provides a linkage between accounts without providing any real assurance of the owner's identity. For higher levels of trust, the use of an easily obtained email account without verification of identity will not suffice.

Trusted identity and its representation online, including protection of individual privacy, were critical issues highlighted in the 2009 White House Cyberspace Policy Review (Executive Office of the President of the U.S. 2009) and Executive Order 13636, entitled *Improving critical infrastructure cybersecurity* (Executive Office of the President of the U.S. 2013). This issue resulted in the NSTIC, which was signed by the President in April 2011 (Executive Office of the President of the U.S. 2011b). According to the White House, there are "many technical and policy shortcomings that have led to insecurity in cyberspace. Among these shortcomings is the online authentication of people and devices" (Executive Office of the President of the U.S. 2011b). The strategy outlined in this document is designed to enhance online choice, efficiency, security, and privacy. Because there may be many different digital identities offered, the interoperability will provide the desired choice of a provider and efficiency in transactions. A program office has been established in the U.S. Department of Commerce at the National Institute of Standards and Technology (NIST) to help move NSTIC forward. Ultimately, this office is intended to be a private-sector-led initiative, and the Identity Ecosystem Steering Group (IDESG) has been established with multiple working groups to see this strategy through to implementation.

One of the goals of NSTIC is that identities be validated in an online environment with minimized disclosure of personal information. Having a multitude of accredited identity providers (IdPs) that are both private and public would offer people a choice of providers. The IdPs will assert information about a user on an as-needed basis. Users could also procure different digital identities from one or more IdPs. For example, a person might want to maintain a professional digital ID, one or more social ones, a health care one, a financial one, and so on. Different IDs will also have different levels of trust or identity proofing. A Facebook- or Gmail-based ID would not have the needed level of identity proofing that might be required for authentication and access to a bank or a health provider application. Multiple IDs also keep a single IdP from having all of an entity's identity information and knowledge of all its interactions. A Private Information Retrieval (PIR) protocol allows users to query large databases while hiding their identity. This type of service could be offered by commercial IdPs and, thus, could enable them to provide some anonymity for their users.

There may be times when anonymous activity within an IT system is desirable both by the users and by those who own and administer the network. For instance, the ability to browse websites or post anonymously can be fundamentally important in some IT systems. The framework under development is very sensitive to this kind of privacy and is ensuring that anonymity is possible when appropriate. The Public Internet Registry, created by the Internet Society (ISOC), supports policy and privacy issues on the Internet, which could include some anonymous communications ("Public interest registry'). There are also efforts in the research community that support limited anonymity capabilities and privacy. For example, the Intelligence Advanced Research Projects Activity (IARPA) started the Security and Privacy Assurance Research (SPAR) Program in 2011. Goals for this program include "implementation of efficient cryptographic protocols for querying a database that keep the query confidential, yet still allow the database owner to determine if the query is authorized" (IARPA). Although anonoymization techniques are generally effective, improved IT capabilities and the increased amount of available data sources could render the originator's identity easier to be discovered.

Hardware

The framework that the IDESG is creating can also apply to Non-Person Entities (NPEs) such as routers, switches, and other hardware. There are other frameworks that rely, in part, on common threat models, but creating such a threat model for supply chain remains a challenge. This is because threat analysis is easier to perform when it is specific to a product and service, and remains context-dependent:

> In a hardware example, an integrated circuit that can be re-programmed after it ships from the original component manufacturer is easier to modify (attack) than an integrated circuit that can only be programmed with a ROM mask during wafer manufacturing. (European Network and Information Security Agency 2012)

The threat is specific and does not apply to software products. In general, hardware threat profiles differ from software threat profiles.

The Trusted Platform Module (TPM) being orchestrated by the Trusted Computing Group (TCG) is another contender for digital identity of hardware. The TPM is a specialized chip integrated into a piece of hardware to securely store digital encryption keys, certificates, and passwords, along with platform measurements that help ensure that the platform remains trustworthy. TPM chips could be used to address supply-chain risk for hardware and could also provide a means to verify the authenticity of the component.

Software

Due to the malleability of software, good digital identity for software requires that software be identified along with attributes to ensure that the software has not changed. Software identification tags (SWIDs) are an emerging technology based on the ISO/IEC standard 19770-2:2009 (ISO/IEC 2009) that facilitates integrity of software packages throughout the supply chain and significantly enhances software asset management. All files (for example, executables, libraries, and scripts) delivered as part of a software product are tagged with the product name from which the file originated, as well as with the vendor, version, hash-code value of the file, and other data. The hash-code value can be verified before the software is installed to ensure that the software has not been altered either accidentally or maliciously in the supply chain, and after installation to ensure that the software remains unaltered. To protect the hash-code value from malicious alteration, the hash-code can be obtained from a trusted third party (for example, NIST), directly from the manufacturer, or encrypted as part of the software package to prevent alteration. This feature allows for definitive software asset management and strengthens the security of the software supply chain considerably.

Software tags provide the capability to identify in real time the status of software residing on systems within an enterprise. This provides a substantial improvement in the knowledge available about the origin and integrity of not only software, but also other associated software elements, such as software patches, service packs, and upgrades. Software tagging would make it significantly harder or nearly impossible to maliciously alter the executables stored on a system and to allow the alterations to remain undetected.

Data

The marking of data objects for identification and the capabilities to manage the data based on those markings would allow increased protection of the data, would improve compliance with corporate and legal policies, and would allow for easy identification of ownership.

One way of marking data objects is through the use of a Universally Unique Identifier (UUID). A UUID is a 16-byte (128-bit) number used by distributed systems to uniquely identify information without significant central coordination (Leach, Mealling & Salz 2005). Thus, any system can create a UUID and be reasonably confident that the same value is not being generated anywhere else in the world. Creating a UUID for each data object creates the foundation for access control, authorization for the acquisition and use of data objects, and the tracking of the movement of data objects. As with software, data is very malleable, and so additional metadata can also be associated with the data object to indicate classification markings, ownership, or any other desired information in a consistent format.

Operational Needs and Challenges

The need for the ability to identify the components of an IT system is not a new problem. In fact, it is a long-standing and difficult problem that extends beyond cyberspace operations and cybersecurity. Companies have struggled for years to precisely identify what hardware is connected to their networks, which software packages are installed on their systems, which files on their systems are related to which software packages, what the current patch level of the software packages is, and how to control data access and flow on their systems. Users continuously struggle with a never-ending list of passwords that must be remembered or, contrary to good security, are either written down or used for multiple accounts. The establishment of valid digital identities for all IT components will not only facilitate better IT management and a better user experience, but it will also help to address these long-standing problems.

The challenges faced in maintaining confidence in IT systems is protecting them from insider threat, from adversary exploitation, and from counterattack. With the implementation of digital identity for IT system entities, IT system users can then be confident that they can access their information as needed and that there has not been unauthorized access or tampering through inadvertent misuse, through malicious insiders, or through external adversaries that would penetrate the enterprise. An effective digital identity capability for all parts of IT systems will ensure that authorized users—anticipated and unanticipated—will have access to the trusted enterprise information and resources they need, when and where they need it, while preventing adversaries' access to the same. However, as these technologies continue to mature and to be deployed, so do the operational challenges. These include the following:

a. Variety: With a growing population of users with increased diversity (such as nationalities, organizational affiliations, operational roles, and security clearances), and an exponential growth of IT devices, establishing a digital identity for each entity is essential. IT systems need to be capable of accommodating unanticipated users where access rules may have to rapidly change in response to political and mission environments.

b. Velocity: Mission tempos are increasing to match cyber speed. This need is a complex one:

> [c]yberspace affords commanders opportunities to make decisions rapidly, conduct operations, and deliver effects at speeds that were previously incomprehensible. However, speed also can degrade cyberspace operations. In some cases a rapid tempo of operations can trigger unintended destruction and evasive actions that would not otherwise have occurred. (Joint Chiefs of Staff 2006)

Consequently, there is a considerable need to ensure that the information from ‘authoritative sources’ remains timely, accurate, and valid.

c. Validity: The U.S. Department of Defense (DoD) is authorized to execute the full range of military operations in and through cyberspace to defeat, dissuade, and deter threats against U.S. interests (Joint Chiefs of Staff 2006). In order to successfully accomplish this mission, it is essential to maintain confidence in national information and to protect it from adversary exploitation. One key component is having identity solutions that are secure and resilient, but are also privacy-enhancing. Additionally, these solutions must also be voluntary to ensure the private sector and international partners are sufficiently incentivized to participate.

d. Volume: The sheer number of entities that comprise an IT system is daunting. The number of software entities on a single piece of hardware can easily be in the hundreds of thousands. The number of pieces of hardware in an IT system for a single organization can be in the thousands or tens of thousands or even greater, as can the number of users. Data numbers can easily surpass the number of software entities.

e. Volatility: By their very nature, IT systems are constantly in a state of flux. New users are added, others removed. Software and hardware are added, removed, updated, or replaced. Data is being added or deleted daily. Such frequent change makes reliable tracking of IT system components nearly impossible using current methods.

f. Verifiability: Having a digital identity for entities in an individual IT system is very important, but many IT systems are connected to other IT systems under the control of another person or entity. Digital identity is the most useful if everyone or everything has it and it has interoperability both inside and outside the enterprise. National security and international cyberspace security cooperation would be enhanced by having a verifiable digital identity capability in place for all IT systems and one that is manually recognized. Among other benefits, this would promote secure sharing of cyber information to respond to cyber incidents (Executive Office of the President of the U.S. 2003). Promoting the need-to-share with mission partners is critical to mission accomplishment. Entities must be able to accommodate federated mechanisms to ensure seamless operation across traditional sovereign boundaries, while retaining protection of sovereign assets for mission partners. Ultimately, IT systems need to be capable of accommodating unanticipated users and other entities where access rules may have to change rapidly in response to political and mission environments.

The Future: Next Steps

The global IT community is experiencing increasing requirements for ease of use of networks and services while maintaining privacy, and the proliferation of a number of 'end points', including sensor clouds, and of new network types, such as vehicular networks. *Identities in the future Internet of Things* (Sarma & Girão 2009) presents new approaches using virtual identities as representations of entities of all kinds as the end points of communications. The increased ease of use and improved flexibility to support new services and means of access in a dynamic and collaborative environment must be matched with an increased ability to quickly identify additions to an IT system and to facilitate the removal of entities that are no longer part of the system. To achieve many of the goals outlined in U.S. cyber policies and strategies, a strong digital identity is needed. Highlighted below are significant technical and policy challenges that will enable the U.S. federal government, as well as other foreign governments, to influence change in the following areas:

<u>Challenge 1</u>: Raise the bar for higher assurance authentication methods for all IT components. The U.S. federal government currently uses NIST SP 800-63-2, *Electronic authentication guideline* (Burr *et al*. 2013), as the process for establishing confidence in user identities electronically presented to an information system. These technical guidelines supplement the U.S. Office of Management and Budget (OMB) guidance, *E-Authentication guidance for federal agencies*, OMB M-04-04 (Office of Management and Budget 2003), which defines four levels of assurance. The OMB guidance defines the required level of authentication assurance in terms of the likely consequences of authentication error. Similar guidelines are needed for all IT components, not just user identities, in order to comprehensively protect an IT system. Authentication levels need to be sufficient to conduct cyber operations and, at the same time, to protect E-commerce applications and to achieve the NSTIC goals.

<u>Challenge 2</u>: Influence standards bodies and vendors to advance digital identities for IT system components. With international participation in the IDESG, there are opportunities to influence the vendor and standards communities by participating in the security, standards, and research working groups. For example, the North American Security Products Organization (NASPO) is developing an American National Standards Institute (ANSI) standard entitled *Requirements and implementation guidelines for assertion, resolution, evidence and verification of personal identity* (ANSI/NASPO-IDPV 2015). This draft standard builds upon the E-Authentication assurance levels and would become an American national standard. Once implemented, it would result in (1) an assurance of identity that specifies an assurance level, (2) a caution or warning message that the asserted identity may not be valid, or (3) referral of the person to exceptions processing. The primary end users of the national identity verification standard will be those entities, both government and commercial, that issue identity credentials people can use to establish eligibility for privileges and benefits. This standard supports the NSTIC's goals and could be used for issuing a digital ID to all U.S. citizens. This standard could also be used by the international community when issuing national ID cards, a protocol which has been adopted by a number of countries.

Challenge 3: Change national policy to require access control for all parts of IT systems.
Executive Order 13587 (Executive Office of the President of the U.S. 2011c) and the *National strategy for information sharing and safeguarding* (Executive Office of the President of the U.S. 2012) are recent publications. Even though these two documents specifically do not address a ‘digital ID’, they do identify the need to improve access control to information. Many of the initiatives in the United States and in the international community support the requirement to improve access control for information. In fact, access control is needed not only for information, but also for all parts of IT systems. Comprehensive access control is a necessary capability to support cybersecurity and cyberspace operations.

Conclusion

Cybersecurity and cyber operations often require considerable resources to implement. Moreover, the components used for cyber operations are unique to and are exclusively used to promote secure information sharing and real-time situational awareness. Better digital identity provides a synergy between cybersecurity and a multitude of other business-related needs, such as asset management, licensing, policy enforcement, disaster recovery, liability, productivity, and even convenience, an attribute that is sometimes considered to be the opposite of security.

The need for comprehensive digital identity is creating the increased demand for new standards; the need for expanding existing standards; the push to develop new architectures, technologies, and new enterprise; as well as the call for shared services to leverage economies of scale and to provide user choice and flexibility while increasing security and supporting cyber operations. The implementation of these demands will pose significant challenges as identity programs and capabilities become more complex, especially because that complexity can increase significantly, if not exponentially.

In order to protect systems against a determined adversary or to prevent a counterattack during a cyber operation, IT system managers need to be able to identify, with precision and in real time, the systems that they have and to ascertain what is on those systems and who is using those systems in an environment that is volatile. If the entities that compromise an IT system cannot be precisely identified, it cannot possibly be known if an IT system has been altered or compromised by an adversary. By relocating trust, verifiable digital identities will move managers and users from a state of not knowing what is on their systems to a state of being able to achieve increasingly effective cybersecurity and cyberspace operations.

References

ANSI/NASPO-IDPV 2015, *Requirements and implementation guidelines for assertion, resolution, evidence, and verification of personal identity*, Document No. NASPO-IDPV-079, version 6.2.

Burr, W, Dodson, D, Newton, E, Perlner, R, Polk, W, Gupta S & Nabbus, E 2013, *Electronic authentication guideline*, NIST Special Publication 800-63-2.

Castro, D 2011, *Explaining international leadership: electronic identification systems*, viewed 22 January 2015, <http://www.itif.org/files/2011-e-id-report.pdf>.

'Digital identity', *Technopedia*, viewed 22 January 2015, <http://www.technopedia.com/definition/23915/digital_identity>.

European Network and Information Security Agency (ENISA) 2012, 'Supply chain integrity', viewed 4 February 2015, <http://www.enisa/europa.eu/activities/identity-and-trust/library/deliverables/sci>.

Executive Office of the President of the U.S. 2003, *The national strategy to secure cyberspace*, <http://georgewbush-whitehouse.archives.gov/pcipb>.
——2009, *Cyberspace policy review: assuring a trusted and resilient information and communications infrastructure*, viewed 22 January 2015, <http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_review_final.pdf>.

——2011a, *Federal Identity, Credential, and Access Management (FICAM) roadmap*, viewed 22 January 2015, <http://www.idmanagement.gov/sites/default/files/documents//FICAM_Roadmap_and_Implementation_Guidance_v2%200_20111202_0.pdf>.

——2011b, *National strategy for trusted identities in cyberspace: enhancing online choice, efficiency, security, and privacy*, viewed 22 January 2015, <http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf>.

——2011c, *Executive Order 13587: Structural reforms to improve the security of classified networks and the responsible sharing and safeguarding of classified information*, 76 Fed. Reg., No. 198, 3 C.F.R. 276.

——2012, *National strategy for information sharing and safeguarding*, viewed 22 January 2015, <http://www.ise.gov/sites/default/files/2012infosharingstrategy.pdf>.

——2013, *Executive Order 13636: Improving critical infrastructure cybersecurity*, 78 Fed. Reg., vol. 33.

Intelligence Advanced Research Project Activity (IARPA) Security and Privacy Assurance Research (SPAR), viewed 22 January 2015, <http://www.iarpa.gov/index.php/research-programs/spar>.

ISO/IEC 2009, *Information Technology–software asset management–Part 2: software identification tag*, ISO/IEC 19770-2:2009, <http://www.iso.org/iso/catalogue_detail.htm?csnumber=53670>.

Joint Chiefs of Staff 2006, *National Military Strategy for Cyberspace Operations (NMS-CO)*, viewed 22 January 2015, <http://www.dod.mil/pubs/foi/joint_staff/jointStaff_jointOperations/07-F-2105doc1.pdf>.

Leach, P, Mealling, M & Salz, R 2005, *A Universally Unique Identifier (UUID) URN namespace*, Internet Engineering Task Force (IETF) RFC 4122.

Office of Management and Budget 2003, *E-Authentication guidance for federal agencies* (OMB Memorandum M-04-04).

'Public interest registry', Internet Society, viewed 21 April 2015, <http://internetsociety.org/public-interest-registry>.

Sarma A & Girão J 2009, *Identities in the future Internet of Things*, viewed 22 January 2015, <http://www.researchgate.net/publication/226218872_Identities_in_the_Future_Internet_of_Things>.

Silver, M, Markus, M & Beath, C 1995, 'The information technology interactive model: a foundation for the MBA core course', *MIS Quarterly*, vol. 19, no. 3, pp. 361-90.

Windley, P 2005, *Digital identity*, O'Reilly Media, Sebastopol, California, United States.

Moving Big-Data Analysis from a 'Forensic Sport' to a 'Contact Sport' Using Machine Learning and Thought Diversity

AJ Ferguson, NM Evans Harris

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

Abstract: *Data characterization, trending, correlation, and sense making are almost always performed after the data is collected. As a result, big-data analysis is an inherently forensic (after-the-fact) process. In order for network defenders to be more effective in the big-data collection, analysis, and intelligence reporting mission space, first-order analysis (initial characterization and correlation) must be a contact sport—that is, must happen at the point and time of contact with the data—on the sensor. This paper will use actionable examples: (1) to advocate for running Machine-Learning (ML) algorithms on the sensor as it will result in more timely, more accurate (fewer false positives), automated, scalable, and usable analyses; (2) discuss why establishing thought-diverse (variety of opinions, perspectives, and positions) analytic teams to perform and produce analysis will not only result in more effective collection, analysis, and sense making, but also increase network defenders' ability to counter and/or neuter adversaries' ability to deny, degrade, and destabilize U.S.* networks.

Keywords: *Thought Diversity, Analysis, Analytics, Machine Learning, Active Cyber Defense, Question-Focused Data Sets*

Introduction

Cyberattacks have evolved in scope, intensity, persistence, and sophistication. This evolution (combined with a commensurate increase in the volume of data these attacks exfiltrate and generate—from gigabytes to terabytes) means the use of *traditional* analytics (for instance, query, combine, filter) will have limited utility and scalability. Specifically, what is needed is analytic tradecraft that allows analysts to query, combine, filter, *characterize,* and *predict* at the 'speed of thought'. The addition of these analytics will allow network defenders to discover non-obvious and unseen patterns in their data. Failing to provide these additional analytics will mean that analysis will continue to be a 'forensic' activity.

A conversation similar to the one below happened at 0900 on June 8, 2014, during a Network Operations Center (NOC) brief to an NOC Director:

> **BRIEFER:** *Ma'am, we know how the compromise happened, who did it, when it happened, the likelihood of its happening again, and the best way to mitigate it. However, it's too late; the destructive malware attack happened four months ago.*
>
> **NOC DIRECTOR:** *Why didn't we catch this attacker BEFORE he could do this damage?*

BRIEFER: *We did not have a signature for the attacker's behavior. We were only able to signature this actor after extensive data collection and analysis. We have no way to reduce our dependency on signatures. While signatures are critically important to our mission, they have a relatively short shelf-life. We looked into using Machine Learning (ML) on the sensor, but the algorithms needed to characterize actor behavior at scale are too complex and computationally intensive for our current sensor infrastructure.*

NOC DIRECTOR*: So why can't we take the analytics we run on the back-end—in the NOC—and put them at the front-end on the sensor?*

BRIEFER*: Ma'am, we could, but our analytics are limited to the capabilities of the product we're using, that is, query, combine, filter. We would have to know what we were looking for* apriori. *We need analytics that can characterize and predict in near-real time and are not computationally prohibitive for our sensor infrastructure.*

NOC DIRECTOR*: So what do we do?*

BRIEFER*: Ma'am, don't know. Hire more analysts?*

NOC DIRECTOR: *Sigh….REALLY!?!*

This conversation happens more often than might be imagined in NOCs around the world—in part, due to the heavy dependency on signatures for detection and subsequent mitigation (Eshel, Moore & Shalev 2014). Optimizing the analysis tradecraft requires the building and deployment of a sensor and sensor infrastructure that enables an effective and timely collision between data and analytics—analysis at the point and time of *contact* (from this point on, referred to as Contact Sport Analysis [CSA]). Just building CSA-enabled sensors is insufficient. The NOCs need thought-diverse analytic teams to perform the level of sense making needed to illuminate adversary tradecraft in new and innovative ways. If NOCs do not embrace thought-diversity-enabled CSA, analysis will continue to be a forensic activity and NOCs will continue to be reactive rather than predictive. This paper makes the following three points:

1. **Investments in Machine Learning will improve cyber analysis** – Deploying CSA will substantially reduce analysis time (because FOA [First-Order Analysis] is performed on the sensor); give more accurate results (fewer false positives); and, because of its behavior emphasis and vice signatures, increase the utility of data brought back for further analysis. In addition, by moving FOA to the sensors, analysts can focus on second-order analysis (trending and root-cause analysis), thereby improving the quality of intelligence products and enabling network defenders to respond, counter, and mitigate more quickly and more effectively.

2. **Thought Diversity is the critical enabler** – Thought Diversity realizes that individuals' thought processes are derived from their unique experiences and, therefore, provide unique perspectives on situations/events. By putting together teams of varying subject matter expertise and analytic approaches, thought-diverse experts can rely upon their intuition and divergent perspectives. As a result, these teams will produce richer, tactical, strategic, and forecast-related intelligence analysis products produced at the 'speed of

thought’. Thought-Diverse ML-enabled analysis will require a paradigm shift that requires analysts to feel comfortable with not analyzing every event the sensor senses.

3. **This shift must occur now** – Resources are scarce everywhere. However, the current analysis process is actually more expensive, in the long run, than a CSA approach.

As shown in **Figure 1** below, one can assume traditional analysis (shown in white arrows) has at least six phases: (1) ingest, (2) extract/transform/load, (3) first-order analysis (FOA), (4) second-order analysis (SOA), (5) intelligence reporting, and (6) counter/mitigate. With CSA, analysts work in a compressed (automated) but more robust process (shown in black arrows) because leveraging Machine Learning on the sensor reduces analysis time while enhancing sense-making insights. Moreover, thought-diverse teams (phase 4 in traditional analysis and phase 2 in CSA) identify non-obvious relationships and patterns sooner rather than later.

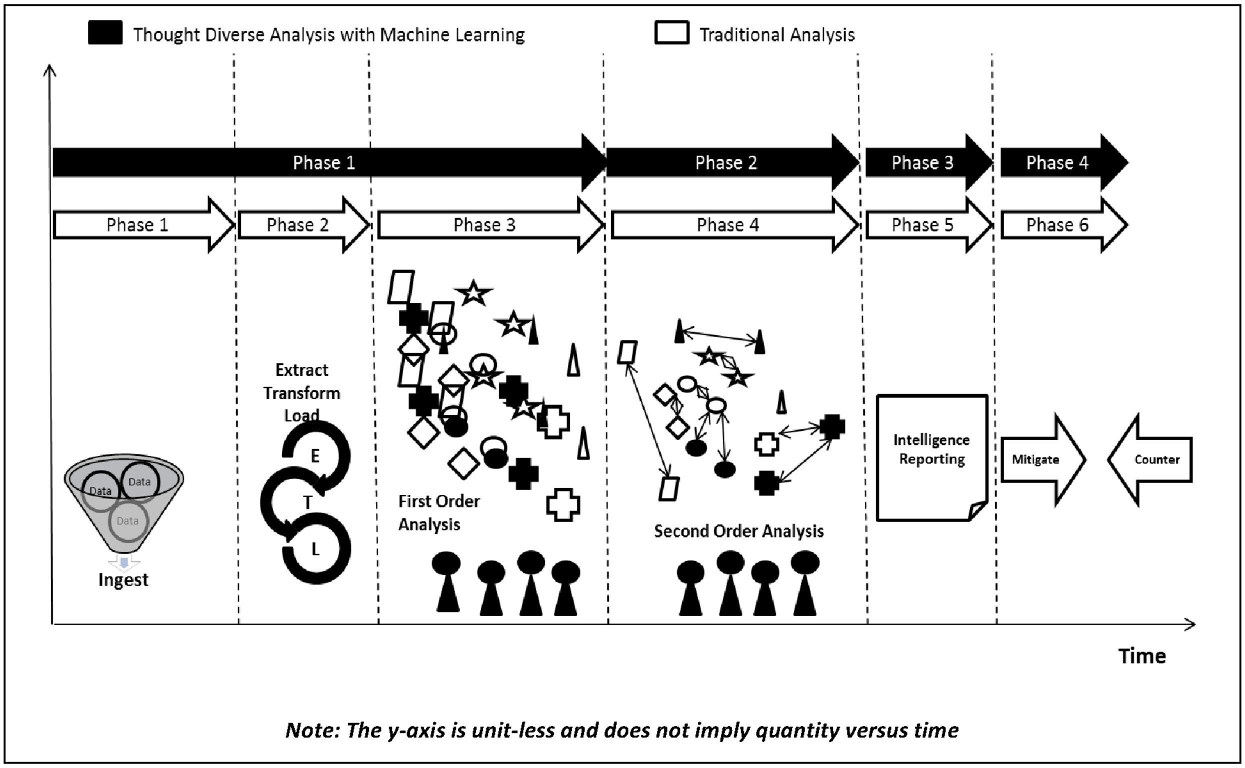

Figure 3: Comparison of Traditional Analysis and Thought-Diverse Analysis with Machine Learning

The next section will describe current analytic challenges and techniques along with their associated benefits and drawbacks.

Current Analytic Challenges

Traditional analysis techniques rely on key words, heuristics, and subject matter expertise for signature development. Even thought-diverse teams that leverage traditional analysis approaches have limited success because they cannot do analysis at scale due to the volume of data. The benefits and drawbacks of these techniques are reflected in **Figure 2** below.

Machine Learning can meet these challenges because, even implemented on current sensors using an ML-enabled analytic platform, ML provides efficiencies and insights while providing increased analytic value (more behavior-based outcomes versus signature-based outcomes). A malware analysis example discussed later will make this point.

Technique	**Benefit**	**Challenge**
Post-collection signature development (heuristics) using thought-homogeneous analytic teams	Allows development of mitigations and countermeasures of a known malicious event or events	When attacks change, new signatures have to be developed, resulting in thousands of signatures without context or knowledge of why certain signatures are ineffective and no insight into the temporal nature of adversary tradecraft.
Thought-diverse analytic teams performing traditional analysis analyze a sample of data using ‘key words’ or ‘business rules’ to reduce the amount of data to be analyzed.	Makes the analysis task more tractable because analysts are analyzing an extract of data that can be an exemplar	Analysts are missing more malicious events than they should because of the volume, variety, and velocity of the data and limited analyst resources.
Thought-diverse Analytic teams using Machine Learning (ML)	Allows teams to work efficiently and effectively because the analysts think differently and have different perspectives	Thought- diverse analytic teams are harder to assemble because they require deep screening.

Figure 2: Current Analytic Techniques, Benefits, and Drawbacks

Another major challenge is cost. There are open-source Machine-Learning technologies available today. For example, Apache Spark, Python SCIKIT-LEARN can be installed, configured, and deployed on existing sensors today, so the increase in cost is negligible. It is important to note that this paper will take the reader down two paths (with examples)—Machine Learning Explained and Thought Diversity. Each will be unpacked separately and then tied together at the end. The next section will briefly explain ML.

Machine Learning in a Nutshell

Machine Learning is the study and application of algorithms that *learn* from data rather than follow explicitly programmed instructions. ML uses supervised and unsupervised learning to discover activity that is similar to something previously seen without having to provide characterization and description information up front (Paxson 2010). Both types are described below.

1. **Supervised Learning (SL):** Analysts apply *apriori*-developed labels, for example, good/bad or member/non-member, to new data instances for the purpose of ‘binning’ events according to their label. Types of SL include:
   - **Classification:** Given samples of malware, the algorithm identifies ‘bins’ samples into two (or more) bins. For example, malicious/not malicious.
   - **Recommender Systems:** These systems produce recommendations by leveraging past behavior as well as similar decisions made by other analysts or through a series of attributes of an object, for example, report author, report theme, or report topic, in order to recommend additional reports with similar attributes.
2. **Unsupervised Learning (UL):** Analysts let relationships in the data emerge automatically, for example, information discovery, *without* the use of *apriori*-developed labels. Types of UL include:
   - **Clustering:** Clustering groups a set of malware samples in such a way that samples in the same group (called a cluster) are more similar to each other than to those in other groups (clusters).

Figure 3 below describes a typical data flow for both supervised and unsupervised learning. It assumes the extract/transform/load process has been performed.

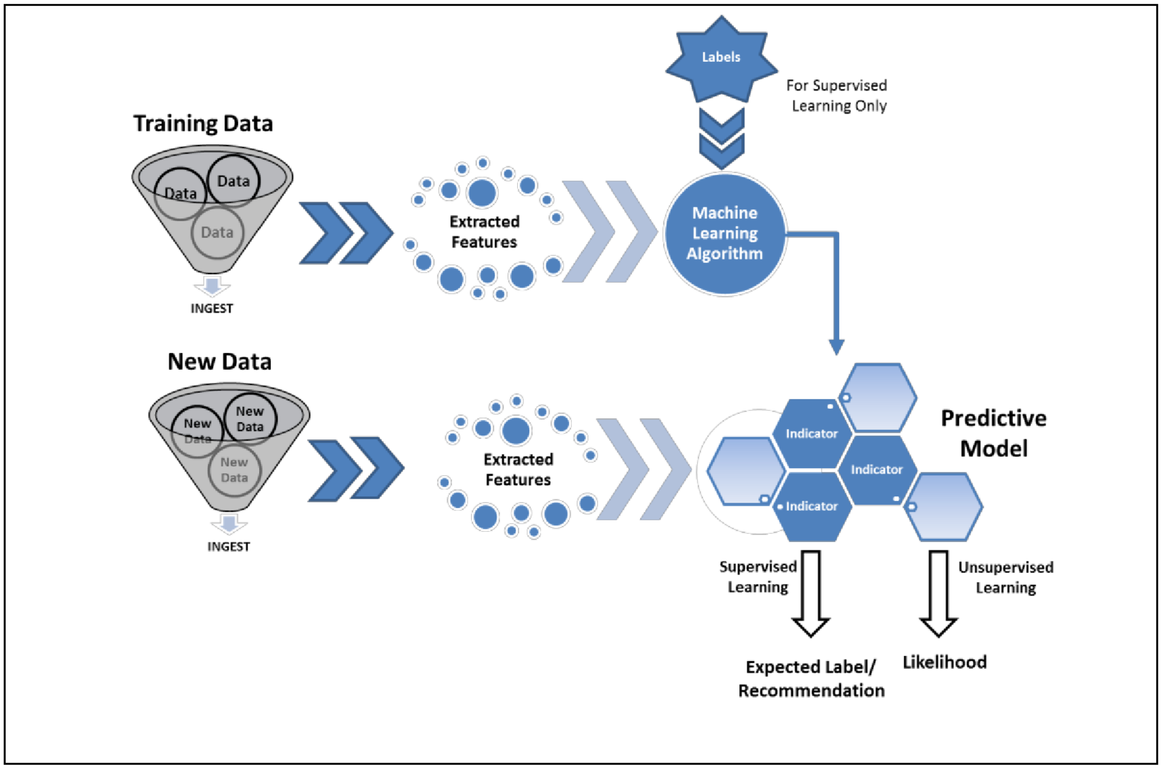

Figure 3: Supervised and Unsupervised Learning Workflow

The Vision: Contact-Sport-Analysis (CSA) Architecture

The authors of this paper propose integrating an In-Memory Cluster Computing (IMCC) module on an existing sensor platform. An IMCC-enabled sensor mitigates several technical challenges. With IMCC, a CSA sensor allows intermediate results to persist in memory, control their partitioning to optimize data storage, and implement iterative map/reduce (cloud) jobs on either a single machine (sensor) or cluster of machines (sensors). For example, Apache Spark allows analysts to run the Machine-Learning libraries on the sensor vice loading larger file systems. **Figure 4** below shows a notional sensor architecture. The sensor may be assumed to have appropriate levels of security to maintain a reliable level of confidentiality, integrity, and availability.

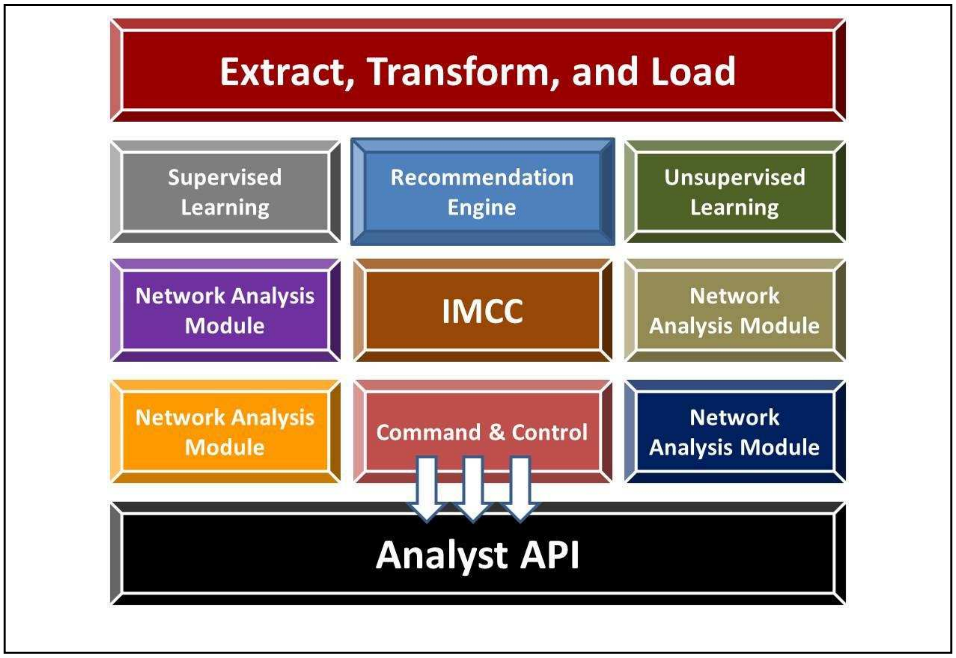

Figure 4: IMCC Machine-Learning-Enabled Sensor

The CSA-enabled sensor will have five components types:

1. **Extract, Transform, and Load (ETL) bus** reads the data, transforms the data into a format that Machine-Learning algorithms can consume, and loads (writes) the data into the target data store.

2. **Three (at a minimum) ML algorithm modules** perform data characterization and include an Unsupervised Learning, Supervised Learning, and Recommender module. Making sure that all three modules have the ability to automatically update predictions (when false positive rates exceed a certain threshold, for example, 5%), and/or adjust

recommendations (when recommenders produce non-intuitive recommendations) is highly recommended.

3. **Network Analysis Modules** can be any reputable network traffic analyzer, for example, Wireshark.

4. **Command and Control (C2) module** controls the proper integration of ML and network-analysis modules, the IMCC, and the ETL bus.

5. **Analyst Application Program Interface (API)** allows analysts to interact with the data.

Even though building a CSA sensor is technically feasible, CSA has its skeptics. These criticisms are discussed below.

Addressing machine-learning criticisms

The four most frequently cited criticisms about using ML algorithms for cyber analysis include the following: (a) the inability to process un-contextualized, unstructured data; (b) their computational intensiveness; (c) their inability to store and retrieve data efficiently from memory; (d); their hard-to-understand and -interpret results. Each criticism is addressed below.

a. **ML algorithms cannot process un-contextualized, unstructured data**: Unstructured data is data that is not stored in a database and has no formal data model. Examples include audio/video files, email, wikis, tweets, presentations, etc. ML can structure and contextualize data through text/content analytics, as well as semantic and statistical analysis.

b. **ML algorithms are computationally intense**: As discussed earlier, IMCC technologies available today can accommodate ML algorithm computational intensity. IMCC enables (a) low-latency computations by caching the working dataset in memory and then performing computations at memory speeds; and (b) efficient iterative algorithms by having subsequent iterations share data through memory, or repeatedly accessing the same dataset. Now developers can easily combine batch, interactive, and streaming jobs in the same application. Once initial models are computed using IMCC technology, the models are run against new data using designated ML algorithms.

c. **ML algorithms need to store and retrieve data efficiently from memory**: IMCC technologies keep track of the data that each analyst produces, and enables applications to reliably store data in memory, essentially allowing applications to avoid costly disk accesses.

d. **ML algorithms results are often hard to interpret and understand**: The returned results are precomputed subsets of data that are derived from the larger data sets and transformed to answer general data characterization questions, also known as Generic Analytic Question-Focused Datasets (GAQFD). These GAQFDs include (1) What is correlated? (2) What is similar? (3) What is good/bad? As the analytic community moves further in to the big-data environment, the need for tools beyond pivot tables for correlation analysis becomes evident.

The ability to run on a single machine, structure unstructured data, mitigate algorithmic complexity, and produce results easy to understand via GAQFDs makes design of a CSA sensor instrumented with ML algorithms the next logical step.

The next section considers Thought Diversity as an underpinning paradigm of CSA. The reader is encouraged to view Thought Diversity as the minimum essential enabler of effective CSA.

Thought Diversity: An Underpinning Paradigm

While ML in a CSA-enabled environment strengthens network defenders' ability to make analytic decisions more quickly, without assembling analytic teams that reflect an adequate amount of thought diversity (no more than five representing at least three different subject matter domains), network defenders still run the risk of deploying signature based upon known behaviors and not truly exploring unknown behavior patterns. Thought Diversity is about realizing that each analyst has a unique blend of identities, cultures, and experiences that define and describe how he or she thinks, interprets, negotiates, and accomplishes a task (Diaz-Uda *et al*. 2014). When forming analytic teams, organizations should consider not only what the analysts know, but also how they approach problems. This is critical to identifying relationships that would traditionally remain unknown and enriching the quality of analytics applied to the ML sensors for scaling across the collection environment (Woods 2008).

With the CSA architecture automating FOA, there is an opportunity to integrate thought diversity into second-order analysis (sense making and root-cause analysis) teams. Having a cyber-intelligence analyst who understands malware activity working with a political scientist who understands open source data and a data scientist who can glean adversary tradecraft based upon data collected can produce behavior-enriched analytics beyond the malware signature. While traditional analysis often brings together different experts, Thought-Diversity analysis brings together different types of expert thinking and perspectives. **Figure 5** below illustrates the future vision for analysis—illumination of the intuitive and counterintuitive.

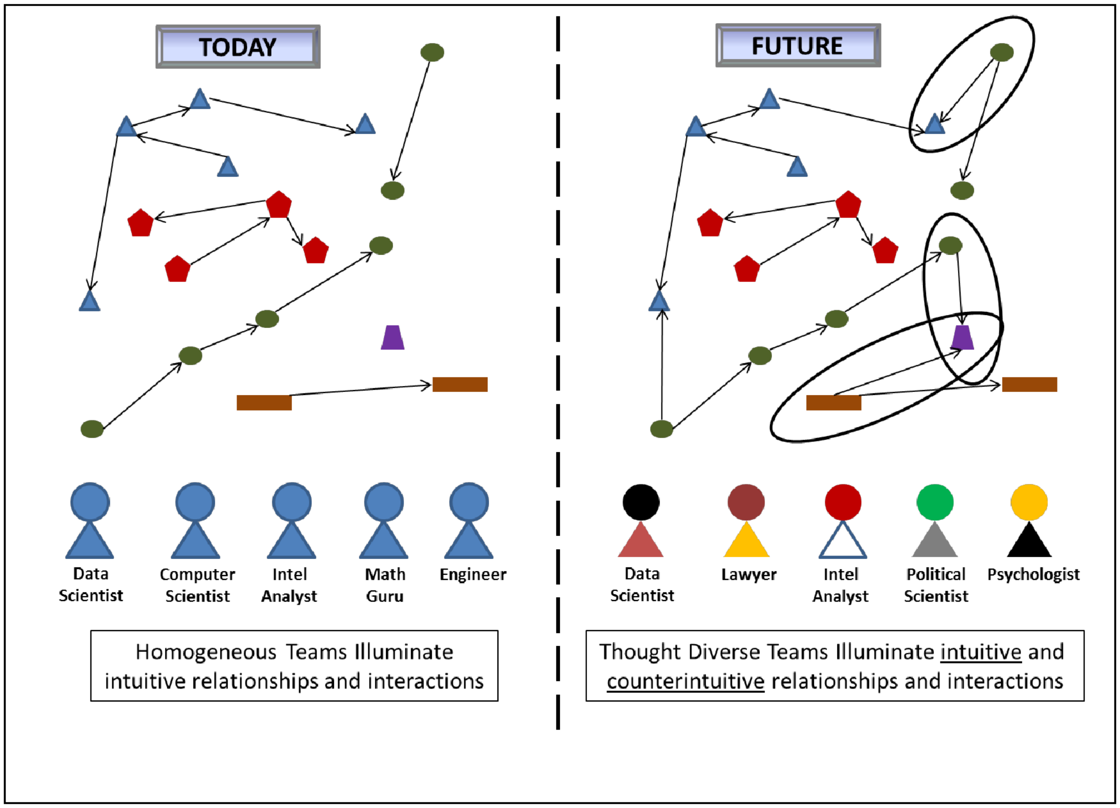
TODAY
FUTURE
Data Scientist
Computer Scientist
Intel Analyst
Math Guru
Engineer
Homogeneous Teams Illuminate intuitive relationships and interactions
Data Scientist
Lawyer
Intel Analyst
Political Scientist
Psychologist
Thought Diverse Teams Illuminate intuitive and counterintuitive relationships and interactions

Figure 5: Traditional (TODAY) Analysis Compared with Thought-Diverse Analysis (FUTURE)

Creating thought diversity: an example

The following example explains one way a thought-diverse team might be assembled. Five candidates were interviewed for two vacant positions on an NOC cyber-analysis team. All candidates were shown the same ML-based analytic output and asked the same 10 questions. Their results are shown in **Figure 6** below.

Many organizations would have assigned Alice, Pablo, and Pierce to a team because they garnered the highest cumulative scores, that is, 50%, 60%, and 70%, respectively. It is safe to assume that Alice, Pablo, and Pierce probably think alike. However, Abegbe's answers reveal an important nuance. While Abegbe was the lowest overall scorer, he correctly answered the two questions the other candidates incorrectly answered—Q1 and Q8. As such, Abegbe presumably brings a different way of thinking to the organization. Hiring Abegbe and Pablo, who approach the questions differently, will increase the team's capacity for creativity. In short, organizations should assemble analytic teams with innovation in mind. This oversimplified example amplifies the need for organizations to recruit people who "challenge the status quo and are self-driven pursuers of their imaginations. The goal is to push, prod, cajole, share, inspire, and enrage as needed to give life to everyone's best ideas" (Hirshberg 1999).

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	% Correct
Alice		**X**		**X**		**X**	**X**			**X**	**50%**
Melanie			**X**	**X**	**X**					**X**	**40%**
Abegbe	**X**					**X**		**X**			**30%**
Pablo		**X**	**X**	**X**	**X**	**X**	**X**				**60%**
Pierce		**X**	**X**	**X**	**X**		**X**		**X**	**X**	**70%**

Figure 6: Analytic Team Candidate Test Results

Machine learning in action: an example

The best way to illustrate the benefits of serious investments in Thought-Diversity-enabled ML on a sensor is by way of a comparison with traditional analysis. This comparison was inspired by Brink (2014).

Alice, the leader of a malware analysis team, gets a few dozen malware samples per week to analyze. The team is able to manually analyze each malware sample and perform the necessary analysis on each sample to decide whether the sample is malicious or not in a few days’ time. Word of this malware analysis team’s proficiency starts to spread, and the number of malware samples submitted for analysis begins to increase—the team is now detecting hundreds of malware samples per week. Although the team tries to stay up with the increased rate of malware samples coming in by working extra hours, the backlog of malware samples continues to grow. **Figure 7** below illustrates this process. The decision-making process takes the malware executable, the metadata, and the behavior history as inputs to make a decision.

As the quantity of malware samples submitted increase, Alice realizes manually processing each malware sample is not scalable. Alice decides to hire and train another analyst, Bob, to help; and Bob allows the team to flush out the malware sample backlog in four days. However, the number of malware analysis requests continues to grow, doubling within a month to 1,000 malware samples detected per week. To keep up with this increased demand, Alice now must hire and train two more analysts. Projecting forward, Alice determines that this pattern of hiring is unsustainable, and she decides to automate some of the decision-making process.

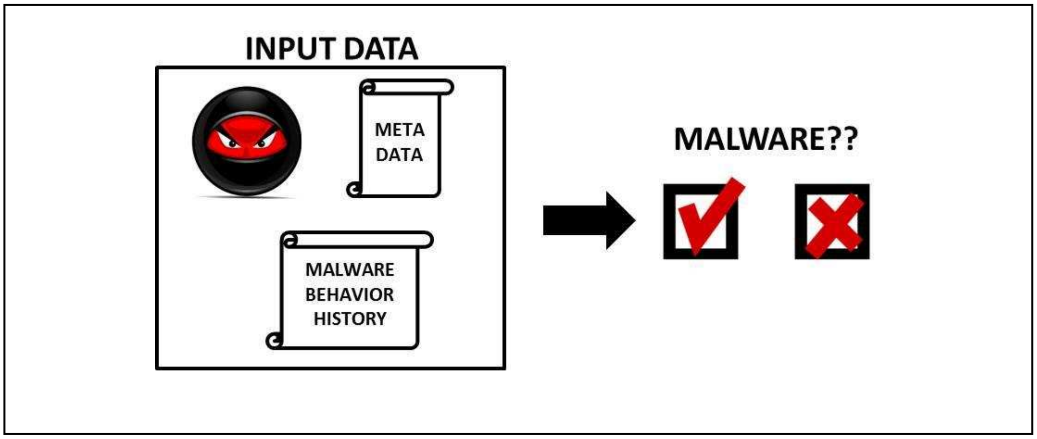

Figure 7: Traditional Malware Analysis Process

After doing malware analysis for several months, Alice’s analysts now have seen the features and behaviors of many of the malware samples that were deemed malicious. Based on her team’s experience, of the 1,500 executables submitted for analysis, 1,000 had a high potential to be malicious. Of the 1,000 deemed potentially malicious, 70% were deemed benign. This set of training data—malware samples labeled as ‘known goods’—is extremely valuable to begin building automation into Alice’s malware analysis process. **Figure 8** below illustrates this process.

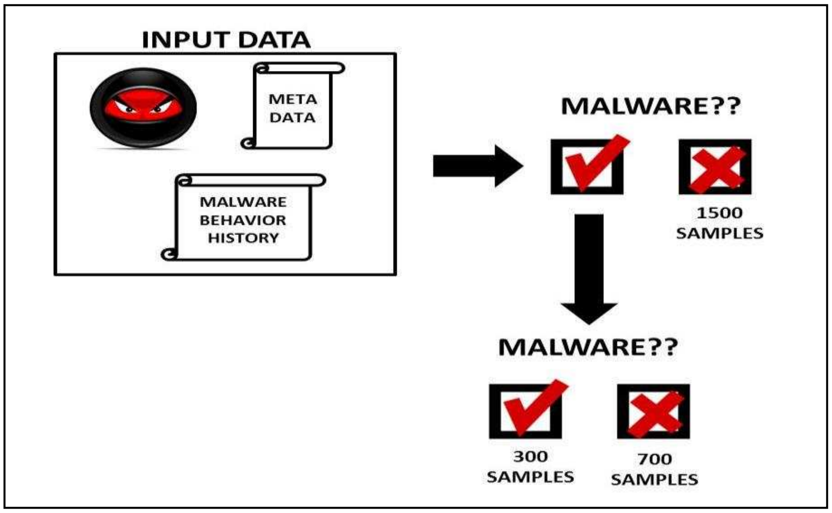

Figure 8: Trending Based on Historical Data

Based upon the scenario above, the traditional sensor versus the Machine-Learning sensor approach can be examined.

Traditional sensor analysis

Using the 1,000 malware samples discussed earlier, Alice's team is now in the position to begin looking for trends. In particular, the team performs a manual search for a set of filtering rules that produce a subset of 'malicious' samples. When new samples come in, the team can quickly filter them through these hard-coded business rules to reduce the number of samples which must be vetted by hand. New trends discovered include the fact that most samples

- use a deterministic beacon time that can be represented as a time series, that is, every hour plus one second, every two hours plus two seconds, etc.
- that use the same IP address for 45 days or less are deemed benign.

Now the team can design a filtering mechanism for their sensor to pare down the number of malware samples they need to process manually through the two rules shown in **Figure 9**.

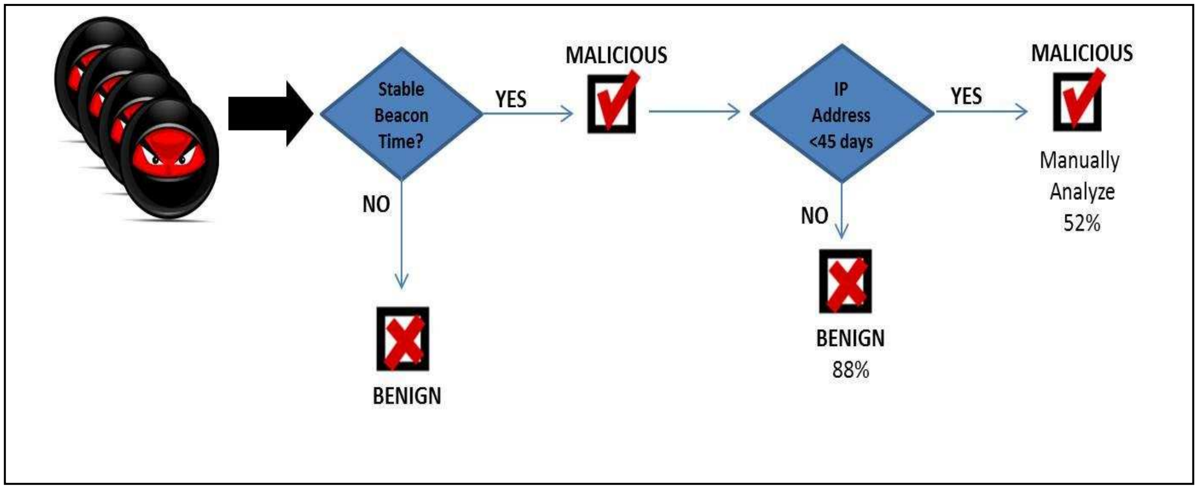

Figure 9: Malware Analysis Implementing Business Rules

The first filter is to automatically label any submitted sample that uses a deterministic beacon time as malicious. Looking through its historical data, the team is able to establish a malicious base rate (MBR) and a Benign Base Rate (BBR). For the MBR, 44 of 86 malware samples turned out to have a deterministic beacon time. So, roughly 51% of these submitted samples were actually malicious. For the BBR, 42 malware samples did not have a deterministic beacon time. This seemed like a great way to quickly identify malware samples. The team then realized that only 8.6% of the submitted samples had deterministic beacon times, which means that the team will still need to manually process more than 90% of the submitted samples. Clearly, the team needs to do some more filtering to get that number down to something more manageable.

The second filter is to automatically label any submitted sample that uses the same IP address for 45 days or less as benign. This seems to be a good filter, as 880 of the 1000 (88%) of the submitted samples used the same IP address for 45 days or less. Including this second filter brings the percentage of malware samples automatically labeled as malicious or benign up to

45% (0.88 multiplied by 0.51). Thus, the team only needs to manually analyze roughly half of the new incoming malware samples.
With these two business rules, the team can now scale its malware-analysis efforts up to twice the amount of volume without having to hire a second analyst, since it now only needs to manually label roughly half of new samples as malicious or benign. Additionally, based on the training set of 1,000 malware samples with known outcome, the team can estimate its filtering mechanism to erroneously label every 70 out of every 1,000 malicious malware samples (7%) as benign and to erroneously label every 70 out of every 1,000 benign malware samples (7%) as malicious. These false-positive rates will be acceptable at first; but as demand for malware analysis grows, the team will have to reduce its false-positive rate.

Under a manual filtering process, there is no way of lowering these error rates without adding additional business rules. This will lead to a number of problems as the volume of malware samples increases: (1) the complexity of signatures will increase making management of them problematic; and (2) there is no statistical rigor to the development of these business rules. Because there is no automated way to do discovery of new/better rules, the team is left to maintain current rules. Moreover, as the patterns of malware sample behavior change over time, the system does not adapt to these changes. The team will have to manually adjust its rules to adapt to changes in malware behavior changes. All of these drawbacks can be traced to the problems that come with using a business-rules approach: the system does not automatically learn from data.

Machine-Learning-enabled sensor analysis

Unlike the business-rules approach discussed earlier, ML performs sense making directly from the data without the need for hard-coded decision rules. Moving from a rules-based to an ML-based decision-making process will provide more accurate and scalable decisions.

To determine whether or not to label a malware sample as malicious or benign, ML utilizes historical training data to predict the best course of action for each new malware sample. To get started with ML for malware analysis, analysts begin by assembling the training data for the 1,000 malware samples. This training data consists of the input data for each malware sample along with the known outcome of whether or not each malware sample was labeled as malicious or benign. The input data in turn consists of a set of features–numerical or categorical metrics that capture the relevant aspects of each malware sample such as the beacon time patterns and IP-address shelf-life. As shown in **Figure 10** below, as new malware samples come in, probabilistic predictions are generated; for example, 83%, 60%, 77%, and 90% of future ‘badness’ or ‘goodness’ are generated instantaneously from the malware sample data. Analysts can group by probabilistic prediction and accept those that meet and exceed a certain threshold, such as 80%.

Machine-Learning models provide more flexible models that can automatically discover complex trends and structure in data without being told what the patterns look like (Brink 2014). ML algorithms, when implemented correctly, can yield a higher accuracy than traditional statistics-based approaches. In this case, that means that the ML model will make fewer mistakes on new malware samples, which translates to fewer false positives.

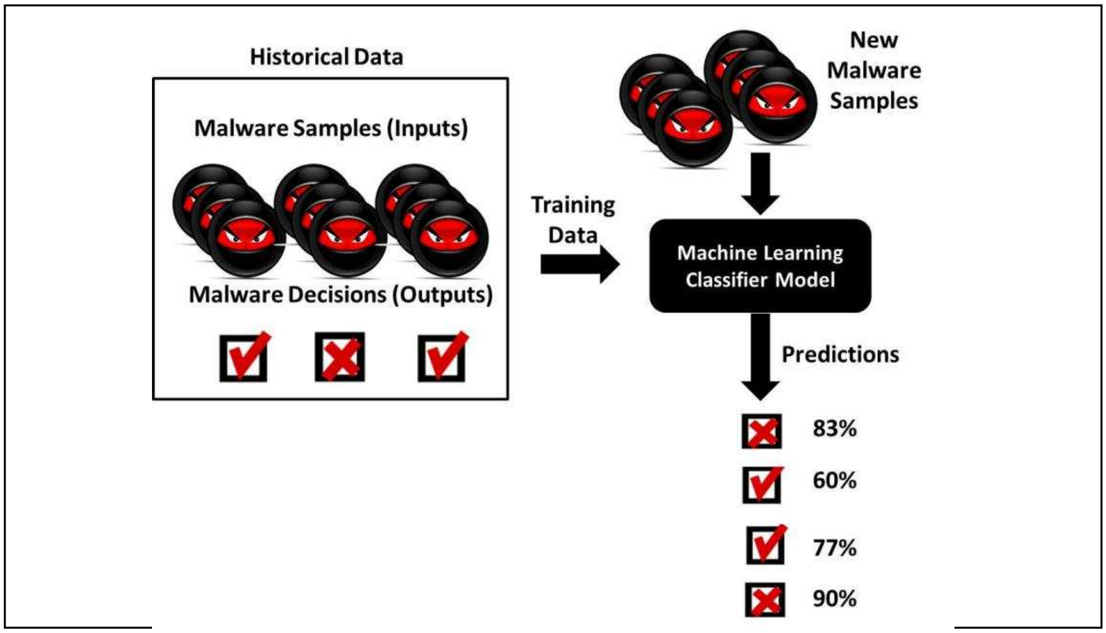

Figure 10: Malware Analysis Using ML Workflow

Figure 11 below shows the number of correctly classified (via ML) Chinese malware samples collected at the beginning of 2015. A statistician might not be able to explain the gap between mid-January and mid-February. However, a political scientist or Intel analyst might notice that this period of time is the Chinese Lunar New Year celebration, so malware attacks would likely decrease.

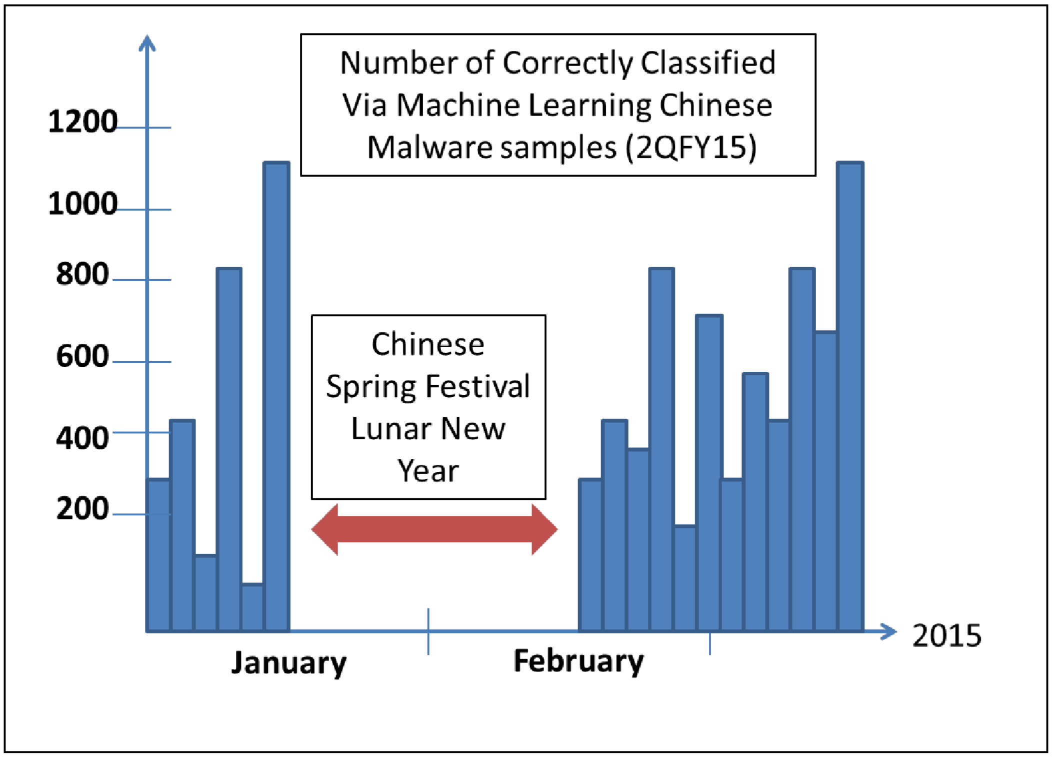

Figure 11: Chinese Malware Samples at the Beginning of 2015

As such, using ML on the sensor with thought-diverse teams has several advantages over manual analysis and hard-coded business rules. They include

1. **Accuracy**: ML uses data to discover the optimal decision-making engine for a given problem. As more data is collected, the accuracy of analysis can increase automatically. Manual analysis and business rules provide a good degree of accuracy until the volume of data increases.

2. **Automation**: As new data comes in, malware samples can be automatically labeled. Hard-coded business rules can be automated, but the business rules have no way to adapt to changes in the data.

3. **Speed**: ML can generate answers in a matter of milliseconds. Manual analysis and hard-coded business rules take time, sometimes hours.

4. **Scalability**: ML easily scales to handle increased data volumes. Most ML processes are parallelizable, enabling infinite scalability. Manual analysis and hard-coded business rules do not.

5. **More nuanced analyses**: ML will help thought-diverse analytic teams illuminate intuitive and counterintuitive correlations and relationships.

6. **Behavior-Based Focus of ML**: The behavior-based focus of ML relieves the need to manage complex signatures.

7. **Statistical rigor**: ML brings statistical rigor to the development of business rules.

The next section will describe a roadmap for implementing a CSA solution.

How to Get Started: A Proposed Solution

As stated earlier, a Machine-Learning-enabled Contact-Sport Sensor can be developed and deployed with existing sensor hardware. The steps are listed below.

1. Instrument an existing sensor with an ETL, Network Analysis, IMCC, C2, ML, and API modules. The Apache Spark Streamer or Python SCIKIT is recommended as the IMCC module.

2. Identify the appropriate MoPs and MoEs. The CSA implementation would have the Measures of Performance and Measures of Effectiveness shown in **Figure 12** below.

3. Next, develop the list of GAQFD questions that need to be answered about the data and describe what format the output should be reported in.

4. Create Thought-Diverse analytic teams by identifying analysts who tend to think differently. For brand-new teams, use an ML-based assessment instrument like the one described in **Figure 12**.

Measures of Performance	**Measures of Effectiveness**
The IMCC Module provides no degradation in overall sensor performance, for example, low latency, CPU utilization, and processor speed.	**Analysis time is reduced by at least 25% while accuracy increases by 25% due to the use of automated machine-learning algorithms. Analysts identify non-obvious relationships between and among data elements that provide insight into adversary tradecraft.**
The IMCC allows analysts to run programs up to 7X faster than Hadoop Map Reduce in memory, or 3X faster on disk.	**False-positive rates for classifier algorithms are less than 10%.**
The IMCC module allows recovery of lost work and operator state without extra coding.	**Analysts can analyze 100% more data due to automated, machine-learning-enabled, first-order analysis capabilities.**
The IMCC allows analysts to simplify their infrastructure by running batch, streaming, and machine learning on the same cluster of data, thus saving thousands of dollars per year.	**Thought-diverse analytic teams are established as evidenced by the variety of insights reflected in reporting.**

Figure 12: Measures of Performance and Effectiveness

5. Have analysts use sensor FOA to inform SOA and recommend mitigations and countermeasures. Identify how the data should be visualized to facilitate second-order analysis.

6. Measure Effectiveness and Performance based on criteria in 2.

7. Refine and adapt.

Why This Is Game Changing

While there are probably dozens of analytic teams deploying Machine Learning, very few (if any) are moving Machine Learning to their network sensors. The approach advocated for in this paper marries technology deployed in new ways used by people with divergent thinking styles to solve an increasingly dynamic and complex set of analytic problems. This approach will result in a measurable improvement in analysis and allow analysis to happen at the 'speed of thought'.

Conclusion

A conversation similar to the following one happened at 0900 on June 8, 2015, during an Intelligence Community Network Operations Center (NOC) brief to the NOC Director:

> ***BRIEFER:*** *Ma'am, we know how the compromise happened, who did it, when it happened, the likelihood of its happening again, and the best way to mitigate it. This attack happened 90 minutes ago, and we predict another one will happen tomorrow at the same time ± 1.5 hrs. We*

are in the process of deploying interim mitigations while we perform additional trending, correlation, and sense making. We moved to a Machine-Learning-enabled sensor that allowed us to characterize the data in near-real time. Now, we are able to characterize this behavior almost immediately. Moreover, we changed the way we formed our analytic teams, so we had competing hypotheses on the table, which broadened the signature to characterize previously unconsidered traits. We were able catch this attacker BEFORE he could do damage!

NOC DIRECTOR: *WOW! Awesome! Have a great weekend!*

References

Brink, H A (n.d.), *Manning Early Access Program (MEAP) real-world Machine Learning, version 4*, viewed August 2014, <http://www.manning.com/brink/RWML_MEAP_CH01.pdf>.

Diaz-Uda, A, Medina, C & Schill, B 2013, ‘Diversity's new frontier: diversity of thought and the future of the workforce’, viewed 5 August 2014, <http://dupress.com/articles/diversitys-new-frontier/#end-notes>.

Eshel, P, Moore, B & Shalev, S 2014, 'Why breach detection is your new must-have, cyber security tool', 6 September, viewed 2 February 2015, <techcrunch.com/2014/09/06/why-breach-detection-ss-your-new-must-have-cyber-security-tool>.

Hirshberg, J 1999, *The creative priority: driving innovative business in the real world*, HarperBusiness.

Woods, S 2008, *Thinking about diversity of thought*, viewed 20 August 2014, <http://www.workforcediversitynetwork.com/docs/Articles/Article_ThinkingAboutDiversityofThought_Woods.pdf>.

Tool References

Apache Spark, <spark.apache.org>.

SCIKIT-LEARN, <scikit-learn/org>.

Wireshark, <wireshark.org>.

On the Role of Malware Analysis for Technical Intelligence in Active Cyber Defense

R Fanelli

Advanced Capabilities Directorate
United States Cyber Command, Fort Meade, Maryland, United States
E-mail: JIWfeedback@nsa.gov

Abstract: *This paper discusses the critical role collection and analysis of malware must play in active cyber defense. The importance of determining the operational characteristics, strengths, and weaknesses of an adversary's weapons and equipment has led to the establishment of technical intelligence (TECHINT) as a discipline in military intelligence. Software, particularly malware, fills the role of weapons in cyberspace. Malware analysis offers significant opportunities to understand adversary capabilities and intent, thus facilitating an effective cyberspace defense. This paper provides background, discusses potential TECHINT gains from malware, and considers how this knowledge may enhance an active cyber-defense strategy.*

Keywords: *Malware Analysis, Active Cyber Defense, Technical Intelligence, TECHINT, Malware Collection, Honeypot*

Introduction

In armed conflict, the importance of determining the operational characteristics, strengths, and weaknesses of an adversary's weapons and equipment has long been understood. This insight has led to the establishment of technical intelligence as a distinct discipline in military intelligence. In the cyberspace domain, software fills the same role of creating effects on targets and facilitating operations that weapons and equipment do in the kinetic domains. However, unlike kinetic-domain weapons, a software weapon can often be rendered completely ineffective by determining its mode of operation and remediating the relevant vulnerabilities and exposures in systems it may target. This potential to negate an adversary's offensive capabilities makes technical intelligence particularly valuable in the cyberspace domain.

Traditional defensive measures, such as software patching, secure configuration, and static perimeter defenses, are necessary, but not sufficient, to defeat sophisticated and persistent attackers. An active cyber-defense strategy must add proactive, intelligence-driven measures to identify, analyze, and mitigate the threats posed by highly capable adversaries.

Journal of Information Warfare (2015) 14.2: 71-85
ISSN 1445-3312 print/ISSN 1445-3347 online

This paper asserts that technical intelligence derived from the collection and analysis of malware is essential to an active cyber defense. The paper starts with a discussion of active cyber defense. It then presents an overview of technical intelligence along with historical examples and examines how this discipline applies in the cyberspace domain. The paper culminates with a discussion of methods for malware collection and analysis, their utility for active cyber defense, and the potential technical intelligence gains from malware analysis before offering concluding remarks.

Active Cyber Defense

What is active cyber defense? There appears to be no single, accepted definition for the term. One definition is “a range of proactive actions that engage the adversary before and during a cyber incident” (Lachow 2013). The U.S. *Department of Defense Strategy for Operations in Cyberspace* offers this definition:

> Active cyber defense is DoD’s synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities. It builds on traditional approaches to defending DoD networks and systems, supplementing best practices with new operating concepts. It operates at network speed by using sensors, software, and intelligence to detect and stop malicious activity before it can affect DoD networks and systems. As intrusions may not always be stopped at the network boundary, DoD will continue to operate and improve upon its advanced sensors to detect, discover, map, and mitigate malicious activity on DoD networks. (2011)

While this definition describes some of what an active cyber defense could achieve, it does not provide much insight into how one might conduct such a defense.

Military doctrine in general contrasts active and passive defense. According to the United States Department of Defense (2010), passive defense entails “measures taken to reduce the probability of and to minimize the effects of damage caused by hostile action without the intentions of taking the initiative”. Conversely, active defense is “employment of limited offensive action and counterattacks to deny a contested area or position to the enemy” (U.S. DoD 2010). Mounting an active defense thus implies the necessity of taking the initiative. Indeed, the description by Lachow (2013) of active cyber defense as “a set of operating concepts that all involve taking the initiative and engaging the adversary in some way” is apt. Active cyber defense does not necessarily require offensive action, but certainly implies seizing the initiative.

Active cyber defense should not be considered synonymous with counterattack or ‘hack back’ (Denning 2008). While the concept of active cyber defense does not necessarily preclude offensive action or counterattack, these are not fundamental elements. Rather, active cyber defense entails an ability to take the initiative and preempt or rapidly mitigate attacks, combined with the intelligence processes to make those actions effective.

A purely passive defense offers a capable and persistent adversary the advantage of a relatively static set of targets and ample time in which to attack. Since operations in the cyberspace domain are largely unaffected by logistical constraints, such as limited fuel or ammunition supplies, the adversary may continue an engagement indefinitely. Given sufficient time and persistence, the attacker is likely to find some weak point in a passive defense. Moreover, a single successful penetration may be sufficient to achieve the attacker’s objectives. For example, an attacker conducted a single, rapid penetration of Stratfor.com in 2011 (Perelstein, Silva & Valentine 2012) that resulted in a significant compromise of sensitive information and damage to the company’s customers and its reputation.

Taking the initiative requires identifying one or more courses of action that will counter an attack in progress or, preferably, prevent the attack from succeeding in the first place. The *intrusion kill chain* concept (Hutchins, Cloppert & Amin 2011) gives a framework for understanding the sequence of actions one must accomplish to engage a target in cyberspace and create the desired effects. This concept is an adaptation of a general *kill chain* framework used in kinetic military operations to describe the series of steps required to identify, localize, and engage a target. Similar frameworks in the field of information system security, such as the Five P’s (*Probe, Penetrate, Persist, Propagate, and Paralyze*) (Cox & Gerg 2004), also relate a sequence of actions an attacker must complete for success.

In active cyber defense, the chain analogy helps to identify potential attacker points of failure that the defender may exploit to defeat the attack. The active defender can take the initiative by identifying links in the intrusion kill chain that can be made cost-prohibitive for the attacker or denied to the attacker entirely, thus reducing or eliminating the probability of attacker success.

The imperative to take the initiative further implies that active cyber defense is an intelligence-driven activity (Hutchins, Cloppert & Amin 2011). The term ‘intelligence’ is used here in the context of military affairs: information about an enemy or the operational environment. Action without effective intelligence is often ill-focused, ineffective, and potentially counterproductive. However, good intelligence, combined with agile command and control processes, facilitates timely, effective action to prevent an attack or mitigate its effects. Thus, implementing an active cyber defense requires attention to effective production and utilization of intelligence.

The field of military intelligence is often divided into disciplines such as human, geospatial, or signals intelligence (U.S. DoD 2007). While multiple intelligence disciplines will contribute to an active cyber defense, a primary enabler is the discipline of technical intelligence.

Technical Intelligence

In the military context, the term ‘technical intelligence’, or TECHINT, refers to the process of collecting and analyzing an actual or potential adversary’s equipment and materiel (U.S. Army 2006). Collection includes items recovered through post-attack forensics, from wrecked or abandoned equipment, from capture during armed conflict, and from acquisition through commercial sources or other third-parties. TECHINT seeks to determine the strengths, weaknesses and other characteristics of the equipment and to assess the technical or scientific capabilities of the adversary. It may also determine methods to render safe adversary munitions and other dangerous materiel. TECHINT helps to prevent technological surprise and often facilitates development of tactics or technology to negate the equipment's strengths or exploit its weaknesses.

TECHINT has historically been an important part of understanding an adversary's capabilities and developing effective responses. Two historical examples from World War II illustrate its value.

The story of the Akutan Zero is a prime example of employing TECHINT against abandoned adversary equipment. In the early part of the war in the Pacific, the Japanese Mitsubishi A6M Zero fighter was dominant. Its speed and maneuverability were superior to the Allied aircraft it faced. Allied pilots avoided dog fighting with the Zero, as they viewed such an engagement as 'certain death.' An Allied aircraft unfortunate enough to have a Zero on its tail was usually doomed (Hammel 1992).

In June 1942, a Zero, damaged in an attack on a U.S. base in the Aleutians Islands, crash-landed nearly intact on Akutan Island (Rearden 1997). U.S. forces subsequently discovered and recovered the Zero. The U.S. repaired the fighter and subjected it to extensive study and flight testing. In addition to determining the aircraft's general capabilities, analysts identified the flight regimes which were most and least advantageous for the Zero relative to U.S. aircraft. For example, the discovery of limitations in the Zero's roll rate and a tendency for its engine to lose power in a negative-G dive led to development of diving and rolling tactics that allowed Allied pilots to reliably escape from a tailing Zero. The TECHINT that the U.S. derived also confirmed that its soon-to-be-deployed F6F Hellcat fighter would be superior to the Zero in nearly every respect.

The Bruenval Raid (Ford 2010), also known as Operation Biting, is an example of an active operation conducted specifically to capture adversary equipment for TECHINT purposes. By late 1941, British Royal Air Force bombers were sustaining a high rate of losses attributed to the detection capabilities of a new German radar system called *Wurzburg*. The British determined it was necessary to capture a copy of the *Wurzburg* in order to develop effective countermeasures. In February 1942, a British airborne raid on a *Wurzburg* site at Bruneval, in occupied France, successfully captured the radar equipment there. The ensuing technical analysis revealed that *Wurzburg* was unaffected by then-current British radar-jamming techniques. The TECHINT gained led to British development of an effective countermeasure in the form of air-dropped, radio-reflective foil strips, or *Window*, tailored to the *Wurzburg's* operating characteristics. *Window* degraded the German radar and significantly reduced the rate of British bomber losses.

As the role of advanced technology in armed conflict has grown, so has the significance of TECHINT. Considering the fundamental role of technology in creating the cyberspace domain and determining its nature, TECHINT is an indispensable source of intelligence there. In the cyberspace domain, software fills many of the roles that weapons and equipment have filled in the traditional, kinetic domains. For example, software serves to locate and gain access to cyberspace targets, and creates effects on those targets.

Collecting and analyzing the software an adversary might use to conduct offensive operations is, therefore, a key element of intelligence gathering and critical to effective active cyber defense. In the realm of information system security, the process of software reverse-engineering and analysis is often referred to simply as *malware analysis*. For convenience, this paper uses the terms *malware* and *malware analysis*, although collection and analysis of any adversary software could yield useful TECHINT.

Malware Analysis and Collection

This section provides an overview of malware analysis, discusses the types of TECHINT gains it can provide for an active cyber defense, and considers malware collection strategies.

Malware analysis involves identifying, characterizing, and understanding the effects and methods of operation of software known or suspected to be malicious. Several texts are available that provide a detailed treatment of this topic (see, for example, Ligh *et al*. 2011; Sikorski & Honig 2012; Skoudis & Zeltser 2003). Malware analysts use a combination of approaches to fully understand the wide variety of malware seen. These approaches can be arranged into three broad categories:

Behavioral analysis involves running a malware sample in a controlled environment to determine its effects. These effects may include file system and network activity, process manipulation, configuration changes, or software persistence.

Static analysis examines the contents of a sample without executing it. This examination may range from simply evaluating any text strings or resources present in a sample to conducting in-depth analysis of program structure or machine code instructions.

Dynamic analysis involves executing the sample in a controlled manner, typically using a software debugging environment to examine and manipulate the sample's internal states. This process facilitates correlating internal events with external effects and achieving more complete analysis, especially for obfuscated or self-modifying samples.

Any non-trivial malware analysis operation requires a *triage* process to prioritize samples available for analysis and focus on those most likely to yield useful results. Malware analysis is a time-intensive activity requiring specialized skills and knowledge, with the quantity of samples generally far exceeding the analytical resources available to examine them. The analytical resources expended on a given sample should be proportional to the useful intelligence likely to be gained. For example, a sample collected from a critical system previously believed to be secure will likely merit more emphasis than, say, a sample collected from a poorly-targeted and unsuccessful phishing attempt.

The triage is often conducted in whole or in part with automated analysis tools capable of rapidly examining large numbers of samples. A variety of tools and techniques for automated malware analysis are available (Egele *et al.* 2012), both commercially (Quinlan 2012; Willems, Holz & Freiling 2007) and as research efforts (Jang, Brumley & Venkataraman 2011; Kirat *et al.* 2013; Raman 2012). Automated analysis tools are most effective in identifying samples of known malware and their variants. However, automated triage can often also identify interesting samples for more detailed human analysis and save time by conducting preliminary analysis tasks and presenting the results to the analyst. Although these automated capabilities are useful, human experts in malware analysis are still necessary. The problem of identifying malware algorithmically in the general case has been shown to be undecidable (Adelman 1992). Malware that is truly novel, highly customized, or resistant to analysis (Branco, Barbosa & Neto 2012) will likely require analysis by a human expert.

Intelligence gains from malware analysis

Malware analysis can produce a wealth of information useful to an active cyber defense. Malware can be viewed as an instantiation of an adversary's tactics, knowledge, technical capabilities, and intentions. This view is particularly true of autonomous malware that must incorporate sufficient knowledge to locate, penetrate, or create effects on a target without interactive human guidance. Static defenses and detection signatures derived only from features of prior samples are brittle in the face of an adaptive adversary using obfuscated, metamorphic, or highly-targeted malware. However, using TECHINT from malware analysis to more completely understand an adversary's goals, attack techniques, and supporting technical infrastructure provides durable knowledge for active cyber defense (MITRE 2012b).

Figure 1 relates examples of the types of intelligence relevant to active cyber defense that may be derived from malware analysis. The examples are organized in the framework of planes (physical, logical, cyber persona, and supervisory) previously introduced to categorize cyberspace operations control features (Fanelli & Conti 2012).

Plane	Example Intelligence Gains
Physical	Hardware identities (serial numbers, models, quantities, or configurations) used for target discrimination (Falliere, Murchu & Chen 2011).
	Vulnerabilities or limitations in hardware that are known to the adversary.
Logical	Vulnerabilities in system and application software of which the adversary is aware. Significantly, this intelligence may reveal previously undisclosed, or *0-day*, vulnerabilities.
	Exploitation methods known to the adversary, including exploits for previously undisclosed vulnerabilities and novel methods to exploit known or suspected vulnerabilities.
	Distinguishing artifacts (for example, unique mutexes, registry values, file or process names) that may signal the presence of the malware on a system (Sikorski & Honig 2012).
	Persistence mechanisms and stealth techniques used by the adversary to maintain access on a target (Sikorski & Honig 2012; Blunden 2009).
	Communications channels and nodes the adversary uses for malware deployment, data exfiltration, or command and control. Examples include network protocols and ports, host addresses, domain names, and generation algorithms (Damballa, Inc. 2012).
	Adversary techniques to obfuscate or encrypt files, network traffic, and other data. This

	intelligence may, for example, reveal cryptographic keys embedded in the malware, custom-built algorithms, or covert channel techniques.
	Programmed behavior of self-propagating malware, such as target search patterns and selection methods.
Cyber Persona	Artifacts supporting attribution of malware authors or users. Examples include file metadata, such as user names, language support, licensing information or EXIF data; artifacts from systems or software used to create the sample; code artifacts such as reused code, distinctive coding methods or defects; and other informative text strings (Hoglund 2010).
	Public Key Infrastructure certificates used for code signing or other authentication.
	Credentials (for example, user ID and password) embedded in the malware.
Supervisory	Command and control capabilities, mechanisms, and command and reply sets.
	Limits on malware propagation and collateral damage (Raymond *et al.* 2013).
	Triggers or timing for initiating and terminating effects.
	Indications of adversary goals and intent. This information may follow from determining the specific information, capabilities, or personas targeted, as well as the effects to be created by the malware.

Figure 1: Malware Analysis Intelligence Gains

Malware, as an instantiation of an adversary’s capabilities and intentions, can provide useful information for an organization committed to conducting the required analysis. However, organizations seeking to mount an active cyber defense must also consider malware sample collection in order to make the most of their analysis capabilities.

Malware sample collection

Collecting relevant malware samples for analysis is a key element of producing useful TECHINT for active cyber defense. In general, it is preferable to collect malware samples earlier and in greater quantity. Collection after the malware has been used can provide useful intelligence to determine the scope of an incident, eradicate any persistent adversary presence, and detect or prevent subsequent attacks. Ideally, however, an organization will collect and analyze samples of an adversary’s most important malware prior to its use in an attack, or at least prior to an attack against the organization. Such proactive collection has the potential to generate timely TECHINT that will enable an active cyber defense to defeat an attack before it occurs.

It bears mentioning that more proactive collection techniques may bring increased resource requirements or issues with legal permissibility. Thus, some organizations will not undertake the full range of collection discussed here.

It is useful to consider three categories of malware sample acquisition: *post-attack collection, post-deployment collection,* and *preemptive collection*.

Post-attack collection refers to collection during or after an incident or attack on the organization. Safeguards such as intrusion detection or prevention systems (IDS/IPS), application-layer gateways and proxy servers, and antivirus systems offer the possibility of collecting malware samples, both from defeated attack attempts and from logged or otherwise preserved samples later found to be malicious. Incident response and forensic processes also offer the possibility of collecting malware samples from persistent media or volatile memory of targeted systems. Similar to the recovery and analysis of the Akutan Zero during World War II, post-attack malware collection and analysis may provide intelligence critical to defeating a persistent adversary or negating the effectiveness of the malware.

This is the least resource-intensive and legally problematic collection category. Any organization seeking to conduct active cyber defense should, at a minimum, have the capability to collect malware samples from day-to-day defensive operations and during incident response. Legal issues are similarly small because the defending organization is acting only within its own systems and is protecting its property.

Post-deployment collection refers to acquisition of samples after they have been released or employed by an adversary, but not subsequent to an attack on the defending organization. Collecting and analyzing malware before an attack can produce TECHINT that facilitates proactive vulnerability mitigation, interdiction of attack or exfiltration vectors, and improved attack detection.

An approach for post-deployment collection is to entice an adversary to act against resources deployed specifically to collect information, such as malware samples, while offering no gain. The general term *honeypot* refers to such an “information system resource whose value lies in unauthorized or illicit use of that resource” (Spitzner 2003). The technique has a relatively long history (Cheswick 1992) and is a mature capability for security practitioners and researchers (Provos & Holz 2007; Spitzner 2002). Honeypots may present a prospective attacker with single or multiple systems, a network of systems (a *honeynet*), or an entire organization. In addition to malware collection, honeypots may also provide warning that an attacker is present, allow for direct observation of an adversary’s interactions with the system, or deceive an adversary into expending time and resources on unproductive attacks. A honeypot may also collect and preserve malware samples that would not be available in post-attack collection due to successful ‘covering of tracks’ by the attacker.

The goal is to induce the adversary to attack a honeypot in lieu of, or at least prior to, attacking production systems. Organizations may deploy these capabilities within their own networks to gain information on attackers specifically targeting them. They may also choose to deploy such collection capabilities in separate locations, perhaps simulating an unrelated organization, to observe a broader range of potential attackers.

Post-deployment collection may also employ a *honeyclient* framework (Gobel & Dewald 2011). This framework uses actual or emulated client-side applications to visit malicious websites or other network services in order to collect samples of malware and information on other client exploitation methods.

Another approach to post-deployment malware sample collection is sharing among organizations. This may take the form of open-source reporting by security researchers and companies, shared malware sample repositories (Contagio 2014; Kernelmode 2014; Offensive Computing 2014), or more closely integrated cyber federations of cooperating organizations (MITRE 2012a). A possibility is for cooperative agreements between organizations with little or no malware analysis capability and those with well-developed capabilities, such as government security establishments, research institutions, or private-sector security vendors. The organizations with analysis capabilities can benefit from a greater, timely inflow of samples while those without benefit from the products of the analysis.

Post-deployment collection will generally require a greater commitment of resources than would post-attack collection alone, especially if the organization chooses to deploy and operate dedicated honeypot systems. Post-deployment collection measures may also be more legally problematic than those used for post-attack collection (Walden 2003). However, this form of collection also offers more opportunity for producing intelligence useful for proactive defensive measures.

Preemptive collection seeks to secure samples before the adversary intentionally deploys the malware. Preemptive collection can gather samples of completed executables as well as proofs of concept, malware source code, or supporting design documents. Since this type of collection permits malware analysis before an attack, it has great potential to produce TECHINT that enables an active cyber defense to proactively defeat the threat.

Observation of adversary communications provides one possible means for preemptive collection. This observation could include capture of network traffic or infiltration of communications means such as Internet Relay Chat (IRC) channels or discussion forums. Such infiltration may involve gaining access via deception or social engineering, by impersonating an existing user or by gaining the cooperation of one or more members of the adversary group. The latter approach has been used successfully by law enforcement to defeat various criminal groups. A notable example is the cooperation of Hector Xavier Monsegur, also known as Sabu, with law enforcement authorities in misdirecting and identifying his fellow LulzSec group members (Bray 2012; Sengupta, 2012). The U.S. Government stated that this cooperation prevented “at least 300 separate computer hacks” and “provided information about vulnerabilities in critical infrastructure” known to the group (U.S. v. Monsegur 2014).

Direct exploitation of adversary-controlled systems provides another avenue for preemptive collection. Systems used for malware development, for command and control, or as malware deployment servers may contain samples that yield valuable TECHINT.

Preemptive collection would, in general, require significant resources and technical capabilities to be successful. While such activity is viewed as permissible in armed conflict and inter-state relations (Humanitarian Policy and Conflict Research [HPCR] 2009; Schmitt 2013), it is likely to be legally problematic in other circumstances (McGee, Sabett & Shah 2013). Thus, this mode of collection not may be available to individuals, businesses, and other private sector organizations. Governmental organizations, however, may have the resources and legal authorities necessary for preemptive malware collection. Moreover, the capacity for government organizations to conduct preemptive collection, organize large-scale post-attack and post-deployment collection, and provide more extensive malware analytic capabilities may be a significant way in which they can contribute to the security of both the public and private sectors.

The primary focus of any offensive component of an active cyber defense should be on preemptive collection of malware samples and other intelligence to assist the defense, rather than on retaliation. Some take the view that active cyber defense implies offensive action, whether as counter-attack or for retaliation (Wong 2011). However, exploitation for intelligence purposes, such as preemptive malware sample collection, promises to be a more effective form of active defense. Intelligence on an adversary's capabilities, intent, and target knowledge can make a direct and lasting contribution to active cyber defense. Conversely, offensive action to disable or destroy adversary-controlled systems and capabilities may offer only temporary relief and runs a much greater risk of collateral damage and other unintended consequences. Here, one of the previously mentioned historical analogies provides illustration: the single British raid at Bruneval provided TECHINT sufficient to defeat German *Wurzburg* radar systems and was a much more effective tactic than attempting to attack and suppress the many *Wurzburg* sites directly.

The collection and analysis of malware must necessarily be a significant element of an active cyber defense strategy. An organization must be prepared to conduct both malware triage and in-depth analysis by human experts in order to mount an intelligence-driven active cyber defense. Organizations must similarly have the means to conduct effective post-attack collection, and, to the extent that resources permit, mechanisms for post-deployment collection. Further, organizations with the requisite legal authorities and resources to conduct pre-emptive malware collection can significantly enhance active cyber defense with the resulting TECHINT.

Conclusion

This paper has explored the role of malware collection and analysis as a critical element of an active cyber defense strategy; it has also considered the nature of active cyber defense and asserted that active cyber defense is about proactive, intelligence-driven defense, rather than counterattack. The author has defined TECHINT as the military intelligence discipline concerned with learning from an adversary's weapons and equipment and asserted that, since software, particularly malware, fills the role of weapons and equipment in the cyberspace domain, the collection and analysis of malware is a primary source of TECHINT in the domain. In addition to providing an overview of malware analysis and presenting examples of the types of actionable intelligence it may provide, the text also defines and discusses categories of malware sample collection, to provide input for an organization's analysis and TECHINT production. Given the fundamental technical nature of the cyberspace domain and the dependence of any proactive defense on intelligence, TECHINT derived from the collection and analysis of malware must be (and must be viewed as) a central element of an active cyber defense strategy.

References

Adelman, L 1988, 'An abstract theory of computer viruses,' *Proceedings of Advances in Cryptology–CRYPTO '88, LNCS 403*, Springer-Verlag, Berlin, Germany.

Blunden, B 2009, *Rootkit arsenal: escape and evasion*, Wordware Publishing, Plano,Texas, United States.

Branco, R, Barbosa G & Neto P 2012, ‘A scientific (but not academic) overview of malware anti-debugging, anti-disassembly and anti-vm technologies’, *Black Hat USA Conference,* Las Vegas NV, viewed 5 January 2015, <http://research.dissect.pe/docs/ blackhat2012-paper.pdf>.

Bray, C 2012, ‘FBI’s ‘Sabu’ hacker was a model informant’, *Wall Street Journal,* 9 March, viewed 5 January 2015, <http://www.wsj.com/>.

Cheswick B 1992, ‘An evening with Berferd, in which a cracker is lured, endured, and studied,’ *Proceedings of the Winter USENIX Conference 1992*, pp. 163-74, viewed 21 February 2015, <http://www.cheswick.com/ches/papers/berferd.pdf>.

Contagio 2014, *Contagio Malware Dump,* viewed 10 January 2015, <http://contagiodump.blogspot.com/>.

Cox, K & Gerg, C 2004, *Managing security with Snort and IDS tools,* O'Reilly Media, Sebastopol, California, United States.

Damballa, Inc. 2012, *Cyber security whitepaper: DGAs in the hands of cyber criminals,* viewed 10 January 2014, <https://www.damballa.com/downloads/r_pub/ WP_DGAs-in-the-Hands-of-Cyber-Criminals.pdf>.

Denning, D 2008, ‘The ethics of cyber conflict,’ *The handbook of information and computer ethics*, eds. K Himma & H Tavani, Wiley, Hoboken, New Jersey, United States.

Egele M, Scholte T, Kirda E & Kruegel C 2012, ‘A survey on automated dynamic malware analysis techniques and tools,’ *ACM computing surveys*, vol. 44, no. 2, pp. 1-49.

Falliere N, Murchu L & Chen E 2011, *W32.Stuxnet dossier,* Symantec Corp., Cupertino, California, United States.

Fanelli R & Conti G 2012, ‘A methodology for cyber operations targeting and control of collateral damage in the context of lawful armed conflict’, *Proceedings of the 4th International Conference on Cyber Conflict*, Tallinn, Estonia, pp. 319-32.

Ford K 2010. *The Bruneval Raid–Operation Biting 1942,* Osprey Publishing, Oxford, United Kingdom.

Gobel J & Dewald A 2011, *Client-honeypots,* R Oldenbourg, Verlag, Munchen.

Hammel E 1992, *Aces against Japan,* Presidio Press, Novato, California, United States.

Hoglund G 2010, ‘Malware attribution: tracking cyber spies and digital criminals,’ *Black Hat USA Conference*, video, viewed 10 January 2015, <http://www.youtube.com/watch? v=k4Ry1trQhDk>.

Humanitarian Policy and Conflict Research (HPCR) 2009, *HPCR manual on international law applicable to air and missile warfare,* The Program on Humanitarian Policy and Conflict Research, Harvard University, Cambridge, Massachusetts, United States.

Hutchins, E, Cloppert, M & Amin, R 2011, 'Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains', *Proceedings of the 6th Annual International Conference on Information Warfare and Security*, Washington, DC, pp. 113-25.

Jang, J, Brumley, D & Venkataraman, S 2011, 'Bitshred: feature hashing malware for scalable triage and semantic analysis,' *Proceedings of the 18th ACM Conference on Computer and Communications Security CCS'11*, ACM, New York, New York, United States, pp. 309-20.

Kernelmode 2014, *Kernelmode.info: a forum for Kernel-mode exploration*, viewed 10 January 2015, <http://www.kernelmode.info/>.

Kirat, D, Nataraj, L, Vigna, G & Manjunath, B 2013, 'SigMal: a static signal processing based malware triage', *Proceedings of the 29th Annual Computer Security Applications Conference, ACSAC '13,* ACM, New York, New York, United States, pp. 89-98.

Lachow, I 2013. *Active cyber defense, a framework for policymakers,* The Center for a New American Security, Washington, DC.

Ligh, M, Adair S, Hartstein B & Richard, M 2011, *Malware analyst's cookbook,* Wiley, Indianapolis, Indiana, United States.

McGee, S, Sabett, R & Shah, A 2013, 'Adequate attribution: a framework for developing a national policy for private sector use of active defense,' *Journal of Business and Technology Law*, vol. 8, no. 1, pp. 1-48.

MITRE 2012a, *A new cyber defense playbook,* MITRE Corp., McLean Virginia, United States.

——2012b, *A public response to emerging exploits*, MITRE Corp., McLean Virginia, United States.

Offensive Computing 2014, *Open malware: community malicious code research and analysis,* viewed 15 December 2014, <http://www.offensivecomputing.net/>.

Perelstein, R, Silva, J & Valentine, J 2012, *Strategic Forecasting, Inc., Computer Forensic Investigation,* Verizon Business Security Solutions report, viewed 10 January 2015, <http://www.scribd.com/doc/229802114/Stratfor-Invesigation-by-Verizon>.

Provos, N & Holz, T 2007, *Virtual honeypots*, Addison-Wesley, Upper Saddle River, New Jersey, United States.

Quinlan, T 2012. *Whitepaper: automated malware analysis,* Norman Shark Co., viewed 25 November 2013, <http://normanshark.com/resource-center/>.

Raman, K 2012, 'Selecting features to classify malware,' *InfoSec Southwest 2012*, Austin, Texas, United States, viewed 21 February 2015, <http://2012.infosecsouthwest.com/files/speaker_materials/ ISSW2012_Selecting_Features_to_Classify_Malware.pdf>.

Raymond, D, Cross, T, Conti, G & Fanelli, R 2013, 'A control measure framework to limit collateral damage and propagation of cyber weapons', *Proceedings of the 5th International Conference on Cyber Conflict,* Tallinn, Estonia, pp. 156-71.

Rearden, J 1997, 'Koga's Zero, an enemy plane that saved American lives,' *Invention and Technology Magazine*, vol. 13, no. 2, viewed 21 February 2015, <http://www.innovationgateway.org/content/koga's-zero-1>.

Schmitt, M, ed. 2013. *The Tallinn manual on the international law applicable to cyber warfare,* Cambridge University Press, New York, United States.

Sengupta, S 2012, 'Arrests sew mistrust inside a clan of hackers', *New York Times,* 7 March, p. A1.

Sikorski, M & Honig, A 2012, *Practical malware analysis,* No Starch Press, San Francisco, California, United States.

Skoudis, E & Zeltser, L 2003, *Malware: fighting malicious code,* Prentice Hall, Upper Saddle River, New Jersey, United States.

Spitzner, L 2002, *Honeypots: tracking hackers,* Addison-Wesley, Upper Saddle River, New Jersey, United States.

——2003, 'Honeypots', *SecurityFocus online forum,* viewed 15 December 2014, <http://www.securityfocus.com/archive/119/>.

U.S. Army 2006, *TECHINT, Multi-service tactics, techniques, and procedures for technical intelligence operations, FM 2-22.401,* U.S. Department of Defense, Washington, DC.

U.S. Department of Defense 2007, *Joint Intelligence, Joint Publication 2-0,* U.S. Department of Defense, Washington, DC.

——2010, *Department of Defense dictionary of military and associated terms, Joint Publication 1-02*, U.S. Department of Defense, Washington, DC.

——2011, *Department of Defense strategy for operations in cyberspace,* U.S. Department of Defense, Washington, DC.

U.S. v. Monsegur 2014, Case 1:11-cr-00666-LAP, Document 30, ‘Sentencing submission by USA as to Hector Xavier Monsegur’, U.S. District Court, Southern District of New York, 23 May 2014.

Walden, I 2003, ‘Honeypots: a sticky legal landscape?’ *Rutgers Computer and Technology Law Journal*, vol. 29, no. 2, pp. 317-70.

Willems, C, Holz, T & Freiling, F 2007, ‘Toward automated dynamic malware analysis using CWSandbox’, *IEEE Security and Privacy*, vol. 5, no. 2, pp. 32-39.

Wong, T 2011, ‘Active cyber defense: enhancing national cyber defense,’ (thesis), Naval Post Graduate School, Monterey, California, United States.

I Want My Smartphone. I Want It Now. And I Want to Connect to Everything from Anywhere…Now!

MLG Althouse

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

Abstract: *Even the classified enterprise is going mobile. Trolls and Luddites cannot prevent it. But the bridge to be crossed to mobility nirvana (a secure, cheap, and user-beloved system) is still rickety with many places where one can fall into the chasm of lost data. The forces of malware, user sloth, shoddy component design, and poor system architecture are arrayed against safe passage. But one is not alone. Assisting the crossing are a number of laws requiring privacy and security measures, government programs that induce superior products, policies written for both public and private sector enterprises, standards bodies, and, most of all customers demanding security from vendors. This paper will look at the mobility mission, the threat to mobile, the secure enterprise architectures using defense in depth, the state of security in system components and how that is being improved through a number of efforts, and the impact of new technology.*

Keywords: *Secure Mobility, Architecture, Policy, Enterprise, Cybersecurity, Risk, Mobile Device*

What the User Wants; What the User Needs

If a capability or feature exists in the consumer mobile space, it is almost certain that there will be some population of enterprise users who will find a business ‘need’ for it. There is generally a significant gulf between what the users can find available on the open market, and thus desire, and the features and capabilities that the enterprise determines are necessary to carry out the mission and are supportable by the enterprise infrastructure.

The Defense Information Systems Agency (DISA) DoD Mobility Enterprise Capabilities lays out the device and system features needed by their customers (DISA). These fall into the following categories: office capabilities, Unified Communication capabilities, collaboration services, enterprise services applications, mission partner applications, device security, and secure access to the Defense Information Systems Network (DISN). Other than the DISN connection and DISA/DoD specific enterprise apps, this list could suffice as an outline for most enterprises.

The Government Mobile and Wireless Security Baseline (Federal CIO Council 2013) provides a broad set of use cases and capabilities for individuals who need to interact with the federal government from the general public to National Security Systems users.

Searching the web for 'enterprise capability requirements for mobile devices' and 'user capability requirements for mobile devices' does not yield as much dispersion as expected. The enterprise results centered strongly on mobile-device management and mobile-app management, as did the user results. Swapping 'user features' for 'user capabilities' only broadened the topics slightly, returning a 2007 study by Gebauer, Yang, and Baimai on user requirements of mobile technology. While out of date in terms of specific technologies, the study's results had functionality characteristics ordered as follows by user survey: multi-functionality, information access, voice, messaging, camera, entertainment, and productivity.

The Enterprise Must Be Protected

Commercial enterprises have a business and fiduciary need, and likely a statutory privacy requirement, to protect enterprise data from loss. U.S. government agencies are regulated by both policy and statute, and a number of these apply specifically to data and network protection.

The Health Insurance Portability and Accountability Act (HIPAA) is probably the most widely known privacy law in the U.S. The Department of Health and Human Services website states,

> The HIPAA Privacy Rule establishes national standards for giving patients the right to access and request amendment of their protected health information (PHI) as well as requesting restrictions on the use or disclosure of such information. The HIPAA Security Rule establishes a national set of security standards for the confidentiality, integrity, and availability of electronic protected health information. (HHS 2015)

Any entity that produces, transmits, processes, or stores PHI is subject to the requirements of HIPAA. This has resulted in security improvements in mobile devices used in the healthcare industry, primarily tablets and laptops, but the technologies, such as encrypted data storage, carry across product lines to consumer devices as well. McLaughlin and Crespo note a number of additional data security regulations from the Federal Drug Administration and Health and Human Services Department, as well as pending legislation to increase protection (2013).

National Institute of Standards and Technology (NIST) Special Publication 800-53 (*Security and privacy controls for federal information systems and organizations*) was created to "provide guidelines for selecting and specifying security controls for organizations and information systems supporting the executive agencies of the federal government" (NIST 2014a). SP800-53 is a foundational IT security requirements document for U.S. government agencies.

Executive Order 13636 (EO13636) *Framework for improving critical infrastructure cybersecurity* issued February 12, 2013, states,

> It is the policy of the United States to enhance the security and resilience of the Nation's critical infrastructure and to maintain a cyber environment that encourages efficiency, innovation, and economic prosperity while promoting safety, security, business confidentiality, privacy, and civil liberties. (Obama 2013)

Critical infrastructures are more numerous than one might think since they include

> systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters. (Obama 2013)

EO13636 tasked the Department of Homeland Security (DHS) to produce a cybersecurity framework, which they have done. *The framework for improving critical infrastructure cybersecurity version 1.0* was published by NIST on 12 February 2014 (NIST 2014b). This dynamic framework document serves as a guide for any enterprise and provides best practices developed by a group of government and private-sector participants.

DHS's National Cybersecurity and Communications Integration Center (NCCIC) is one entity that spans all domains: defense, intelligence, civil critical infrastructures, and law enforcement/counterintelligence (DHS 2015). The U.S. Computer Emergency Readiness Team (US_CERT) falls under NCCIC, as does the equivalent organization for industrial control systems.

The Department of Defense (DoD) cybersecurity is governed by a number of documents, beginning with DoD Instruction 8500.01 (Takai 2014). It addresses cyberspace defense, integration and interoperability, identity assurance, cybersecurity workforce, and risk management. The DoD CIO Terry Takai also created the Department of Defense Commercial Mobile Device Implementation Plan to specifically guide the department's programs to deploy commercial mobile technology (Takai 2013). The two main thrusts are that Mobile Device Management (MDM) should be part of any mobile device deployment and that the devices obtain their apps from an enterprise Mobile Application Store. It also addresses FirstNet, the national cellular first-responder network that DoD personnel can use. FirstNet should provide the enterprise greater insight into how its devices are behaving on the cellular transport part of the system.

What's the Worst That Could Happen?

This technology is cool; what could go wrong? Next, this paper examines the threats to mobile devices and what might result if one of these threat actors conducts an attack.

What are the threats? This analysis begins by assessing what a person has that someone else might want or might benefit from having. The threat comes from the person or organization that wants the information or property that the target has. The threat actors below are categorized by their motivation.

- Criminals: If it has value and can be converted to money there will be someone who will be interested in taking the information or data. These are the inveterate phishers, dangling various baits at the end of an email, web page, or text message. They invite readers to just click on this tasty morsel to be rewarded. Only later will the pull of the line be felt. Many of the big news commercial data breaches are the result of criminal organizations stealing the personal and financial information of employees and customers. There is a flourishing international market for this data. Ransomware is a burgeoning and profitable criminal pursuit that will be discussed later.

- Amateur hackers/enthusiasts: These people are out to gain street credibility, are curious, or are just practicing skills and techniques.

- Industrial espionage: This is criminal activity but is practiced by a different class of actors seeking information of value to a commercial client. It is all about intellectual property.

- Terrorists: The goal could be destruction, publicity, money, or leverage to negotiate for some other outcome.

- Political ideology: While it has several facets, the Sony hack falls mainly into this category. The North Korean attackers intended to effect a policy change at Sony (Do not release the movie *The Interview*!) through publication of embarrassing email, posting of movies not yet in theaters (financial loss), and threat of further information releases.

- Nation-State espionage: These are the top-shelf adversaries, although nation states comprise a spectrum of ability. And their targets comprise a spectrum as well. Governments want information about the intentions, capabilities, finances, and personnel of other governments or entities. Nation states also tend to have the largest budgets, the best trained operatives, access to needed technology, and persistence. National policy generally dictates the targets and policy changes slowly. Thus, a nation state can often take a more lengthy and careful attack path because the goal is enduring.

What Are the Bad Outcomes?

If an attacker is successful, what can he or she achieve? These are the events that computer and network defense are intended to prevent. Damage to an enterprise has to be measured individually. What is worse: having a contact list stolen, losing an email, missing an appointment, having a website defaced, or suffering a false financial transaction? It depends. The following bad outcomes are ordered loosely by the increasing sophistication of the attack needed to accomplish them.

Denial of service to the user can be a rather crude attack that can result in annoyance or in significant loss. Malware can corrupt or disable a device. There are many bits of code, apps, or even malformed packets that can place a device into an unstable state. The worst of them will break a device such that it cannot be recovered, and might only be able to be reset to the factory default state. Any attack that can disrupt the normal operation of the device has the potential to prevent a user from carrying out his or her required business functions.

Enterprise security measures detect an attack on or intrusion of the device and deny the user enterprise access. This is a good thing from the perspective of the enterprise, but the user still suffers. If an attacker’s goal is to take mobile employees offline, then deploying some noisy cheap attack that triggers the device's or enterprise's defenses to deny the device access could suffice.

Ransomware denies a user access to his or her data. Both individuals and enterprises are victims that have to pay the attacker for the means to decrypt the victim's data which the attacker has encrypted. Open source encryption algorithms available to attackers, such as AES, are strong enough that a victim's only options are to pay up or to abandon the encrypted data.

If the data on the device itself is not protected by encryption, then it is quite straightforward to remove it if an attacker gains physical access. Unencrypted data is more likely to be accessible by a variety of applications including any malware. If a password or particular app is needed to access the data, it places another barrier in front of the attacker. An attacker could also remotely access encrypted data files, but then he or she would need some means to decrypt them to gain access.

The data on the device can become corrupted and made unrecoverable. Even if the attack does not result in data theft, the data can be lost to the user. Malware can corrupt the data on the device intentionally, as a result of trying to access it, or as a byproduct of corrupting the storage, memory, or I/O systems. When the attacker is not paid, ransomware is a form of this attack.

User credentials (username, password, and certificate) are removed from the device and are used to access the enterprise from an attacker's system. If loss of the user credentials is not detected or reported, the attacker can masquerade as the user and gain 'legitimate' access to the enterprise and all the information that the user is authorized to see.

Enterprise data can have considerable value. It could be customer information for financial gain (credit card or banking information), client lists, intellectual property (designs, plans, blueprints, proposals, etc.), or employee data. General examples of enterprise intrusions are Sony (though there were significant misdeeds beyond data theft in this case), Target, Home Depot, JP Morgan Chase, and Staples (Hardekopf 2015). In none of these cases was a mobile device known to be an attack vector.

Enterprise data is corrupted, and trust in data integrity is lost. This could be a worse outcome than theft of the data. Imagine if the enterprise's financial data is corrupted (changed) and not detected. The corruption propagates through the backup system until there is no unaltered copy remaining. Now the enterprise is faced with financial data that cannot be trusted. The attacker could also insert email, orders, instructions, or events into the enterprise business systems that might sow discord and disruption.

Use of a mobile device to infiltrate the enterprise is one of Bitdefender's Top 5 threat predictions for 2015 (Ban 2015). A mobile device can be used to introduce malware to enterprise systems. Once a device is compromised, the attacker can leverage the user's legitimate enterprise access to survey the parts of the enterprise network accessible by the user and to look for vulnerabilities or weak internal security measures that can be exploited. If the user has permission to install applications, then the attacker can take over the user's internal machine or can add capability to a roaming profile.

The mobile device can be used as an exfiltration path. A legitimate user with access to the enterprise opens a communication channel between the mobile device and the enterprise that can be leveraged by malware already inside the enterprise network to route data through the mobile device's wireless public network interfaces to the attacker's network. Because the enterprise has some trust in the mobile device communication routed to it, the information may come under less scrutiny and thus may be a less risky way to extract information. Using WiFi or Bluetooth, the attacker could avoid moving data across the Internet by using a listening post to communicate with the victim's mobile device at convenient times and locations, such as at home, at a regular stop, or during a commute.

Secure Mobility Architectures

A mobility architecture must first support the enterprise services and capabilities required by the user. Making it secure, especially for classified networks, traditionally meant development of purpose-built devices and software that was largely stand-alone in terms of security. This became an unsustainable path due to the cost and development time; the solutions were out of date by the time they were ready to deploy.

The National Security Agency (NSA) developed a secure mobile system architecture that is described in the Mobile Access Capability Package (MACP) (NSA 2014). It is based on commercial components and standard protocols. The MACP is derived from the first CP in the Commercial Solutions for Classified (CSfC) Program, the Mobility Capability Package, and is the basis for several system instantiations accredited for classified use (NSA 2015b). The MACP calls for the following fundamental security measures.

Encryption of data in transit (DiT) is a basic privacy measure and is one of the oldest security measures in networking. The MACP calls for two independent encryption layers using standard protocols and Suite B algorithms (NSA 2015c).

A layer of DiT that provides an additional security function is a Virtual Private Network or VPN. The VPN is an encrypted tunnel between two devices—in this case between the mobile device and a VPN concentrator at the boundary of the enterprise. Having an always-on VPN means that the mobile device's IP traffics all routes to the enterprise where the traffic can be inspected. If the VPN is not always-on, then the mobile device can communicate directly with the Internet, thus exposing it to attackers. This is called 'split tunneling'.

User and device certificates should be stored in hardware-backed memory, such as a Secure Element, Trusted Platform Module, or Hardware Security Module. These chips implement cryptographic keying functions that interact with applications while keeping keys and certificates secure. The certificates can be the most valuable thing on a mobile device, since they can provide an attacker valid access to the enterprise networks and data.

Mobile device management (MDM) systems typically consist of a client app on the mobile device and a server in the enterprise. The MDM client enforces a security policy set by the enterprise, but this is done at the application layer of the mobile device, and, thus, the MDM has less control permission. The MDM is also limited by the application programming interfaces

(APIs) that the manufacturer makes available to the application layer. This varies widely between manufacturers and requires MDM vendors to produce multiple versions of the client app tailored for each device model. Devices with a smaller market share may get few or no compatible MDM clients. Enterprise Mobility Management (EMM) is a newer term that encompasses MDMs and adds in system, user, and enterprise services policies.

Monitoring is one of the key security components of a composed commercial solution. The solution is viewed as an entire system and not just the user device; monitoring must be done wherever and whenever possible. Intrusions into the system can occur at almost any point. The defense of the system is based on the layered security—defense in depth. It is a given that there are vulnerabilities in the components and weaknesses in the architecture. The layering of security measures is designed to require that the attacker defeat multiple barriers before gaining access to data. Wherever these barriers are, they should be monitored for failure. The MACP has a new network domain, generally referred to as the Grey network. In a traditional classified network, there is a Black or unclassified domain and a Red or classified domain. The connection between these domains can be a Type 1 certified encryption device, a guard, or a cross-domain solution. Classified data is protected by a single layer of strong encryption when transiting an unclassified network resulting in a Black/Red interface point. In a CSfC architecture, there are at least two layers of commercial encryption. They are terminated on the enterprise side by two discrete devices; thus, there are two transition points and an intermediate state between Black and Red, which is Grey. The Grey domain contains no user or enterprise application processes, just singly encrypted traffic transiting it. The management plane should be out of band. Areas between defensive layers, such as the Grey domain, are ideal places to monitor for intrusions. The traffic and flows in the Grey should be well behaved and follow just a few protocols. It is a quiet space as opposed to the noisy Black and Red and amenable to a small tight rule set. An attacker must cross the Grey to reach the Red and diverge from the normal flows to do so.

Encryption of Data at Rest (DaR) is necessary if sensitive data will be stored on the mobile device. Devices will be lost and stolen. Recovering the user's and enterprise's data from an unencrypted device is straightforward. Tools for this job abound, and free software from the device manufacturer designed to let the user manage the device will often suffice. Even with DaR enabled, there are many poor implementations. If the device has no hardware-backed memory for key storage, the encryption key will be written to the system partition of the device's memory and can be recovered with forensic tools.

A hardware Root of Trust (RoT) is the security anchor of a mobile device. It is an immutable function that is designed into the chipset. When the device boots, low-level hardware functions load firmware into memory. The firmware performs most of the basic computing functions, and it also loads the operating system. An RoT authenticates the firmware and attests to its integrity. With the firmware now trusted, it can perform an integrity check on the operating system, and the operating system in turn checks applications. Thus, the device can start from a known good state. This is not a perfect defense because the system can be dynamically corrupted once it begins to communicate with external devices. But the user knows that, as long as the device has not been physically tampered with, a reboot will return it to a known good state.

Another desktop and server security mechanism now appearing in the mobile space is the Virtual Mobile Device (VMD). A VMD in the form of a virtual machine mobile operating system is launched on a server in the enterprise. An app on the physical device connects to the server and the VMD supplants the display of the physical device with that of the VMD as a thin client. To the user, there is no difference in the operating environments. The VMD's apps run on the server and, most importantly, all enterprise data stays in the enterprise. An enterprise Bring Your Own Device (BYOD) capability can be enabled by placing one VMD client app on the mobile device. If the device is agency owned, then more device security should be applied with an MDM and other host-based measures. The user certificate can stay in the enterprise where it is more secure, and a device certificate can be used for authentication with the enterprise boundary.

Enterprise-side security measures include encryption (TLS 1.2) of the connection between the device and the enterprise, a VPN layer for additional security and risk reduction if desired, a boundary firewall, isolation between the VMDs within the server, process controls (MAC policy) in the server to limit resource access by the VMDs, and management of VMDs by an MDM client.

Determining the security value or strength of a product can be a daunting task for even a well-staffed enterprise. The primary sources of information are product literature (naturally biased), reviews by magazines, security websites and bloggers, and word of mouth from colleagues. An unbiased evaluation has been difficult to come by. The CSfC program has a listing of approved component products that can be used to compose a secure solution following the architecture in a Capability Package (NSA 2015a). Mobile technology is the fastest growing area in this list. The process to become a listed CSfC component starts with certification by a National Information Assurance Partnership (NIAP 2015) accredited lab for compliance with a NIAP Protection Profile. Protection Profiles are a set of security requirements, both threshold and optional, and include assurance activities the lab uses to validate the security functions. There is a whole family of mobility-related Protection Profiles covering all component products that perform a security function (NIAP 2015). While not every component used in a CP has a listed product yet, as this list grows, it will become a valued resource for secure system designers.

Mobile Technology Future

The near-term mobile technology future holds some promise to deliver products with improved security postures and some new mechanisms that can be added to existing secure architectures to provide additional layers of defense.

Virtualization creates a non-physical instance of a device that is hosted by a physical device. Primarily this has been within servers, but workstations and other types of devices can run a virtual machine (VM) as well. The interface between the physical device and the VM is called a hypervisor. There are two primary types: Type 1 sits between the VM and the hardware while Type 2 sits between the VM and host operating system (Popek & Goldberg 1974). These VMs have a level of isolation or indirection from the physical device and its operating system. According to one study,

> That indirection makes it possible to draw a neat line around everything that is inside a virtual machine and clearly distinguish it from what is on the outside. By carefully

> managing the interface between the virtual machine and its environment, a hypervisor can prevent software running in a virtual machine from directly accessing the hypervisor, other virtual machines, or the rest of the outside world. (Sandia 2011)

Wearables might become tokens for the mobile devices providing a physically separate store for credentials, keys, and data. They would communicate via Bluetooth or Near Field Communication (NFC). The concept is that the wearable and mobile device are paired and must be sufficiently close, within some specified distance, for the mobile device to access the wearable and be able to function in a secure mode. Nick Jones (2014) cautions that "Personal accessories such as smart watches displaying email and messages will pose new security and management challenges for employers. Devices that can record video will raise many privacy concerns, as has been demonstrated by Google Glass" (Jones 2014).

A recent post by Emma Ban (2015) of Bitdefender on the *Top 5 security predictions in terms of technology developments and practices in 2015* listed machine learning algorithms for attack detection and network defense, and stronger BYOD policies as numbers one and two on the list. Given the rapidity with which malware can morph to defeat signature-based detection schemes, behavior-based techniques that are self-learning about the device they protect seem to hold great promise.

Can Secure Mobility Be Achieved?

What is secure enough? It is a risk decision. The inputs to the risk 'equation' are vulnerabilities, threat, security measures and mitigations in the system, the benefit of the system to the enterprise's ability to carry out its mission, the cost of an attack, and the value to the attacker if an attack is successful. There is no accepted mathematical equation to tie all of these variables together.

Once a system, mobile or traditional, is built, someone must make the decision to deploy it. In the U.S. government, this person is the Accreditation Official (AO), and he or she is often also the CIO or CISO. The 'go' or 'no go' decision is based on risk. CSfC compliant systems have a risk analysis performed as part of the development and integration process. A risk analysis is produced as part of a Capability Package, but this is based on the architecture alone. Known vulnerabilities in the components may not be fully mitigated and may not have a single defensive layer protecting them. When a system is instantiated, the specific products used should be NIAP Certified and CSfC listed; if so, a body of evaluation data exists for them. Other component products require some individual evaluation. Analysts then look at various attack paths to determine the conditions and likelihood of success and the cost of executing the attack. There are vulnerabilities in every device and system. At least one must be exploitable, and there is a cost to exploiting it: discovery cost, exploit-development cost, delivery cost, and the potential cost of being exposed should the exploit be discovered and attributed to the attacker. Risks are then categorized, often as high, medium, or low. A more granular and measurable process is desirable.

Insurance companies are starting to offer cyber-risk insurance, so professional associations (such as the Risk Management Society) are developing industry processes to assess cyber risk. The U.S. government also has a Risk Management Framework (RMF) that transforms the prior Certification & Accreditation processes (RMF 2014). The RMF was developed to ensure compliance with policy and is implemented through a number of NIST and FIPS publications as noted in **Figure 1**, below (NIST 2014a). A circular process such as this should be followed and repeated on a regular basis because the operating environment (and, thus, the premises under which the risk decision was made) is continually changing.

Secure mobility can be achieved. It is a matter of an individual enterprise's decision based on the risk and reward. Some enterprises have said yes and others no, and there is a spectrum of mobile capabilities and accesses that have been deployed based on risk acceptance. Commercial mobile technology is becoming more secure because of privacy laws, the move to perform banking and financial transactions with mobile devices, and customer demand from both consumers and enterprises. This positive development is conditioned by a parallel increase in threat actor focus on mobile technology.

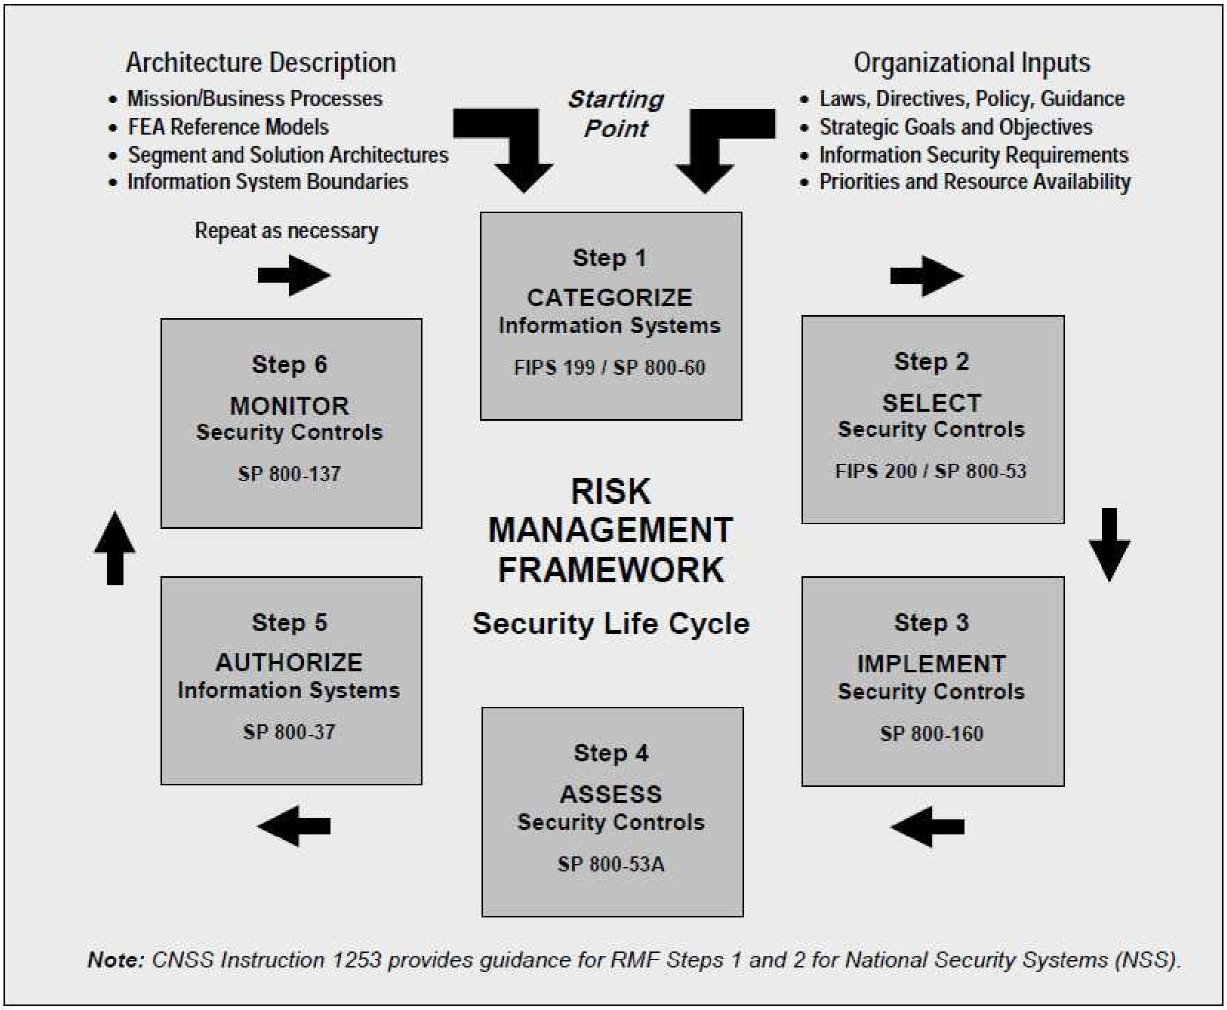
Architecture Description
• Mission/Business Processes
• FEA Reference Models
• Segment and Solution Architectures
• Information System Boundaries
Starting Point
Organizational Inputs
• Laws, Directives, Policy, Guidance
• Strategic Goals and Objectives
• Information Security Requirements
• Priorities and Resource Availability
Repeat as necessary
Step 1 CATEGORIZE Information Systems FIPS 199 / SP 800-60
Step 2 SELECT Security Controls FIPS 200 / SP 800-53
Step 3 IMPLEMENT Security Controls SP 800-160
Step 4 ASSESS Security Controls SP 800-53A
Step 5 AUTHORIZE Information Systems SP 800-37
Step 6 MONITOR Security Controls SP 800-137
RISK MANAGEMENT FRAMEWORK
Security Life Cycle
Note: CNSS Instruction 1253 provides guidance for RMF Steps 1 and 2 for National Security Systems (NSS).

Figure 1: Risk Management Framework

Conclusion

Mobile devices are here to stay, and secure integration of them into the workplace poses significant issues. The technology itself is largely consumer-driven and is transforming at a faster pace than most enterprises can accommodate and faster than security policy can be developed. Security capability itself in mobile devices follows behind development of new features since the shiniest 'dooflatchey' is what drives the market. Thankfully, consumer demand for privacy and regulation surrounding banking and payment apps have shortened the lag between a new device or application's release and implementation of security measures. The security professional faced with protecting an enterprise connected to mobile devices will always be challenged to determine the risk of permitting use of the latest mobility breakthrough and finding sufficient mitigating mechanisms to bring the risk into balance with the reward. The employment future of the cyber defender is bright.

References

Ban, E 2015, *Top OEM predictions for 2015: security becomes smarter*, viewed 17 January 2015, <http://oemhub.bitdefender.com/top-oem-predictions-for-2015-security-becomes-smarter>.

Department of Homeland Security (DHS) 2015, *US CERT*, viewed 7 Jan. 2015, <https://www.us-cert.gov/nccic>.

Defense Information Systems Agency (DISA) 2014, *Enterprise capabilities*, viewed 12 December 2014, <http://www.disa.mil/Services/Enterprise-Services/Mobility/Enterprise-Capabilities >.

Federal CIO Council 2013, 'Government mobile and wireless security baseline', viewed 25 January 2015, <https://cio.gov/wp-content/uploads/downloads/2013/05/Federal-Mobile-Security-Baseline.pdf>.

Gebauer, J, Tang, Y & Baimai, C 2008, 'User requirements of mobile technology − results from a content analysis of user reviews', *Information Systems and e-Business Management* vol. 6, no. 4, pp. 361-84.

Hardekopf B 2015, 'The big data breaches of 2014', *Forbes*, viewed 17 January 2015, <http://www.forbes.com/sites/moneybuilder/2015/01/13/the-big-data-breaches-of-2014/>.

Health and Human Services (HHS) 2015, 'Your mobile device and health information privacy and security', viewed 14 February 2015, <http://www.healthit.gov/providers-professionals/your-mobile-device-and-health-information-privacy-and-security>.

Jones N 2014, 'Top 10 mobile technologies and capabilities for 2015 and 2016', Gartner Report G00260239.

McLaughlin P & Crespo M 2013, 'The proliferation of mobile devices and apps for health care: promises and risks', Bloomberg BNA, viewed 13 February 2015, <http://www.bna.com/the-proliferation-of-mobile-devices-and-apps-for-health-care-promises-and-risks/>.

National Information Assurance Partnership (NIAP) 2015, *Protection profiles*, viewed 7 January 2015, <https://www.niap-ccevs.org/pp/>.

National Institute of Standards and Technology (NIST) 2014a, *SP 800-53 Rev. 4: NIST Special Publication 800-53 Revision 4, Security and privacy controls for federal information systems and organizations*, viewed 13 January 2015 <http://dx.doi.org/10.6028/NIST.SP.800-53r4>.

——2014b, *Framework for improving critical infrastructure cybersecurity V1.0*, National Institute of Standards and Technology, viewed 23 January 2015, <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214-final.pdf>.

National Security Agency (NSA) 2014, *Mobility security guide*, viewed 27 December 2014, <https://www.nsa.gov/ia/_files/Mobility_Security_Guide.pdf>.

———2015a, *CSfC component list*, viewed 4 January 2015, <https://www.nsa.gov/ia/programs/csfc_program/component_list.shtml >.

———2015b, *CSfC program*, viewed 4 January 2015, <https://www.nsa.gov/ia/programs/csfc_program/index.shtml>.

———2015c, *Suite B cryptography*, viewed 4 January 2015, <https://www.nsa.gov/ia/programs/suiteb_cryptography/>.

Obama, Barack, *Presidential executive order 13636*, viewed 12 December 2014, <http://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity>.

Popek, G & Goldberg, R 1974, 'Formal requirements for virtualizable third generation architectures', *Communications of the ACM*, vol. 17, no. 7, pp. 412–21.

Risk Management Framework 2014, *Risk Management Framework*, viewed 27 December 2014 <http://www.rmf.org/index.php/what-is-rmf.html>.

Sandia National Laboratories 2011, 'The security impact of system virtualization', *The Next Wave*, vol. 18, no. 4, <https://www.nsa.gov/research/tnw/tnw184/articles/pdfs/TNW_18_4_Web.pdf>.

Takai T 2013, *Department of Defense commercial mobile device implementation plan*, viewed 25 January 2015, <http://www.defense.gov/news/dodcMdimplementationplan.pdf >.

——2014, 'Department of Defense instruction 8500.01', viewed 23 January 2015 <http://www.dtic.mil/whs/directives/corres/pdf/850001_2014.pdf>.

Defending Cyberspace with Software-Defined Networks

GH Bishop, SR Boyer, MJ Buhler, AJ Gerthoffer, BC Larish

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

Abstract: *Software-Defined networking (SDN) presents a new way of thinking about and operating communication networks that is revolutionizing the networking industry. This paper first describes how a core tenet of SDN—a logically centralized network control plane—enables dynamic, fast, and predictable changes in network behavior. Next, the authors show how network operators can use this capability to transform defensive cyber operations from today's labor-intensive, static processes into automated, agile responses that are capable of dealing with tomorrow's cyber threats.*

Keywords: *Software-Defined Networks (SDN), OpenFlow, Defensive Cyber Operations (DCO)*

Introduction

Over the last several years, the computer networking community has begun to think differently about designing and operating communications networks. In the past, network designers and operators focused on individual boxes and the network protocols running on those boxes. Recent attention has turned to a more holistic view of the network whereby a centralized mechanism allows designers and operators to treat the network as a single entity. This new perspective, called Software-Defined Networking or Software-Defined Networks (SDN), represents a fundamental shift in thinking and presents a significant opportunity for new types of Defensive Cyber Operations (DCO).

Because SDN is such a fundamental change and a relatively new idea, many groups have developed their own (sometimes conflicting) definitions of the concept. Although these definitions will likely converge as the technology matures, for the purposes of this paper, the authors propose a definition based on Open Networking Foundation (2015b). SDN is a communications network exhibiting two characteristics:

1. Modular network hardware and software and
2. Logically centralized control using OpenFlow as a building block.

The first SDN characteristic can be described by drawing a comparison with traditional network devices. Traditional devices are sold by vendors that bundle hardware with their own proprietary software. In contrast, SDN allows consumers to combine network hardware from one vendor with network software from another vendor. This model is similar to a computer on which consumers can install one vendor's operating system on a different vendor's hardware.

The second SDN characteristic requires a more in-depth description of a network device. Typically, network devices are described as consisting of a number of 'planes' that implement different functionality. For example, a device's data plane takes packets that arrive on an input port and forwards them out of an output port, and a device's control plane instructs the data plane how to forward different packets out of the device's various ports. (For example, a router chooses an output port based on a packet's destination IP addresses.) In a traditional network device, the control and data planes are bundled together inside the device, and they communicate with each other via the vendor's proprietary protocol. In contrast, an SDN separates the control plane from the data plane and standardizes OpenFlow as the protocol they use to communicate. This model allows a network device's control plane to exist in a physically separate location from its data plane. As a result, the control plane can be implemented on multiple physical servers that the network device views as one logical server (which is often called an SDN Controller), and that logical server can act as the control plane for multiple devices in the network (Open Network Operating System 2015; OpenDaylight 2015; Ryu 2015). Furthermore, applications can be written on top of the SDN controller to make the network behave in certain ways.

Figure 1, below, illustrates this architecture. This situation is analogous to how a supervisory module controls multiple line cards in a chassis router. The difference is that SDN enables the physical separation of the supervisory module functionality from the line cards, and standardizes the protocol for communication. A key consequence for DCO is that network devices will only forward packets into the network if the SDN Controller has explicitly instructed the device how to handle those packets. Although many alternatives to OpenFlow have been proposed, the authors focus on OpenFlow because, at the time of this writing, it is the only multi-vendor protocol that gives the control plane fine-grained control over the data plane. That fine-grained control enables many of the DCO capabilities discussed below.

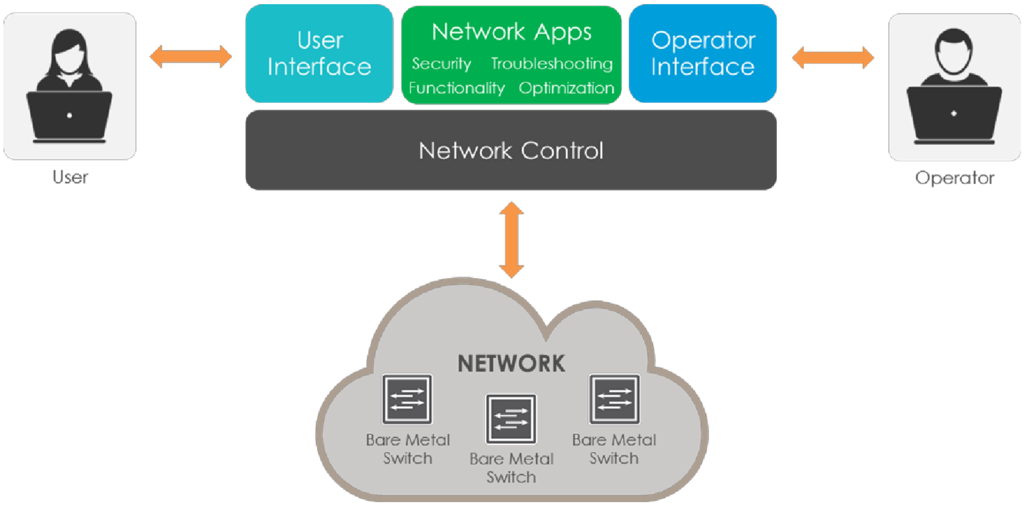

Figure 1: SDN Architecture

Proponents of SDN argue that these two characteristics will provide numerous benefits to network operators, including reduced Capital Expenditures (CAPEX), reduced Operational Expenditures (OPEX), and more innovative network capabilities (Heller *et al*. 2013). Because the primary benefit of modular network hardware and software is reduced CAPEX and OPEX, the rest of this paper will not focus on that aspect of SDN. In contrast, logically centralized control of the network via OpenFlow can not only lower costs, but it can also create opportunities to implement capabilities that improve defensive cyber operations. For example, centralized network control allows an organization to explicitly specify what and how devices connect to the network; fine-grained control means that network operators can monitor and respond to very specific data flows; and centralized control combined with fine-grained control results in predictable network performance that operators can use to respond to network events. Many of these ideas appear in the first research efforts related to SDN (Casado *et al*. 2007).

The rest of this paper describes SDN capabilities that improve DCO (removing learning from the network, dynamic access control, and rapid response), describes open questions related to SDN and DCO, and concludes with a review of the paper and a look toward the future

Eliminating Network Learning

One reason for the Internet's great success is the ease with which different organizations can connect via a communications network (Caesar *et al*. 2010). Traditional network learning protocols, such as Spanning-Tree Protocol (STP), Open-Shortest-Path-First (OSPF) routing, and Address-Resolution Protocol (ARP), result in a 'default-allow' connectivity model in which networks dynamically and automatically enable communications with any new devices added to the network (Comer 2013). However, as organizations became more dependent on communications networks for mission-critical functions, the default-allow connectivity model was no longer acceptable; organizations needed the ability to restrict which devices were allowed to communicate in order to guarantee network performance, network availability, data separation, etc. As a result, mechanisms such as Virtual Local Area Networks (VLANs) and Access Control Lists (ACLs) were layered on top of the default-allow connectivity model, in effect creating a 'default-allow-restrict-later' model. From a DCO perspective, this model is equivalent to building in security after the fact, which often results in a patchwork of rules that makes it difficult to implement security policies or even determine whether desired security policies are actually being enforced. Furthermore, this model opens up attack vectors for spoofing and cache poisoning. Instead of layering additional rules on top of existing learning protocols, the authors propose using SDN to create a secure network by eliminating the learning protocols entirely.

A description of how to eliminate one learning protocol, ARP, in a Data Center Network (DCN) using SDN's logically centralized control illustrates how this model might work. This model begins with the premise that implementing a secure network requires knowing what devices are connected to the network. Next, the model leverages the fact that modern data centers use automation tools (for example, Puppet, Chef) to ensure servers in the data center maintain a known configuration. In particular, these tools are used to store network topology information, as well as to preconfigure the IP addresses of servers according to the servers' physical locations in the data center.

Since a network device in an SDN will not forward packets into the network until the controller has provided the device with instructions on how to do so, an application can be written on top of the SDN Controller to read the topology, MAC address, and IP address information from the automation tools and program the network devices as follows.

First, network devices are instructed to forward packets out of particular ports based on the packets' destination IP addresses. If the device is sending a packet directly to a server, the device is also instructed to ensure that the destination MAC address of the packet matches the MAC address of the server by rewriting the packet's destination MAC address. (The reason for this will become clear in the next paragraph.) Proper forwarding behavior is guaranteed because the controller has received complete topology and IP address information for the DCN from the automation tools.

Second, network devices are instructed to send all ARP packets to a specific server called the ARP Monitor. (For illustrative purposes, an application on top of the SDN Controller can be assumed to act as the ARP Monitor.) The ARP Monitor is configured to generate an ARP reply containing a 'dummy' MAC address (for example, 00:00:00:00:00:00) any time the Monitor receives an ARP request. Correct forwarding behavior is still guaranteed because, as described above, packets in the network are forwarded based on IP address. Thus, the MAC address in an ARP reply is irrelevant; the only issue of significance at this point is that the server transmits the packet on the network.

Operating in this manner facilitates DCO for the DCN in a number of ways. To begin with, it reduces the attack surface of the system by eliminating unauthorized ARP packets from the network; since servers will only ever receive ARP packets from the ARP Monitor, spoofing and cache-poisoning attacks are not possible. It also allows for a more comprehensive monitoring solution because all ARP packets originating from the network will be sent to the ARP Monitor, and the ARP Monitor can be configured to alert on unexpected packets. If, for instance, a device in the middle of the network receives an ARP reply, operators know there is a problem because ARP replies should only be sent from the ARP Monitor to a server. While this problem could be as simple as a switch's failing or a port going down, the ARP reply could also indicate that a host has joined the network at an unauthorized location. In either case, the SDN can generate an alert that something is wrong and can allow the network to take action. Finally, this model can monitor all packets in the network to ensure they contain the 'dummy' destination MAC address. If not, the ARP mechanism has somehow been subverted, and appropriate action can be taken.

Three scalability issues must be addressed for this approach to be practical. First, changes in network topology (for example, addition or removal of a network device or server) must be accounted for. As stated above, this model assumes that any secure network must know the devices connected to it. Therefore, changes in network topology are addressed by requiring that the network be static enough that all device additions and removals can be recorded in the data center automation tools. Changes due to temporary link failures are addressed using redundancy in the network, as suggested by Caesar, Casado, Koponen, Rexford, and Shenker (2010). The second scalability issue is related to the amount of forwarding information the controller must configure on each network device. This issue is addressed by picking an IP addressing scheme to reduce the number of forwarding rules the network device needs. Lastly, sending all ARP packets can generate a high load on the controller in large networks. This load can be significantly reduced by using the ‘dummy’ MAC address since it eliminates the need to do a lookup on the MAC address to IP address mapping.

Although the current discussion focusses on a data-center implementation, this model applies to any network whose configuration changes can be recorded in a database or automation tool. In other words, this model can be applied to any network whose configuration changes on a timescale of hours to days. Campus networks, for instance, are a possibility. Also worth noting is that similar approaches can be taken to eliminate other learning protocols, such as STP and routing.

Dynamic Access Control

In today’s networks, dynamic access control is difficult to perform because of the limited data available to operators and the limited control those operators have over network devices (Nayak *et al*. 2009). To collect data that might be used to make access control decisions (for example, log-in information or packets a device generates), operators generally must deploy some type of hardware or software sensor. To take action on the information the sensor generates, operators will typically reconfigure network devices manually. Automated solutions exist, but they often involve proprietary hardware and software, or they require software that can only support a few select pieces of networking hardware. This mode of operation is not scalable and will not be effective against the cyber threats of the future.

By allowing fine-grained control over how network devices forward packets, SDN overcomes today’s limitations and gives network operators the ability to control network access in a truly dynamic manner. The following description of how passive fingerprinting could be used to control an end-point’s network access illustrates this advantage.

First, the OpenFlow protocol provides fine-grained control over a network device’s traffic-forwarding behavior (Open Networking Foundation 2015b). An SDN Controller instructs a network device how to forward traffic by sending it a list of rules that dictate how various packets should be handled. An OpenFlow rule contains two elements: a match and an action. The match element of a rule lists characteristics the network device should use to determine if a particular rule applies to a packet (for example, the Layer 2 – Layer 4 header fields). The action element of a rule tells the network device what to do with a packet if a particular rule applies. As a specific example, an OpenFlow rule might instruct a network device to match packets with a particular destination MAC address and take the action of forwarding the packet out of a particular device port.

The remainder of this section describes how this capability can be used to take advantage of the well-known fact that the Dynamic-Host Configuration Protocol (DHCP) can be utilized to passively fingerprint a device’s operating system in a typical scenario, and how SDN’s flexibility allows for a number of different implementations.

An application on top of the SDN Controller instructs all network devices to send all DHCP packets they receive to a DHCP Monitor. (For illustrative purposes, it may be assumed the application also acts as the DHCP Monitor.) The DHCP Monitor then generates a fingerprint for the device based on the DHCP packet and compares the fingerprint to a whitelist containing known good fingerprints. If the fingerprint is on the whitelist, the DHCP Monitor sends the DHCP packet back to the network device along with instructions to forward that packet to the DHCP server. If the fingerprint is not on the whitelist, the SDN Controller instructs the network device to physically shut down the port where the DHCP packet arrived.

The flexibility provided by the OpenFlow protocol allows network operators to implement this in a variety of ways that are not possible with traditional networks. For example, instead of requiring that the DHCP Monitor check fingerprints before allowing access, the SDN Controller can instruct the network device to forward DHCP packets to the DHCP server as normal and also send a copy of the DHCP packet to the monitor. In this case, access is granted initially, but it can be revoked if the monitor's offline analysis returns a negative result. Since the DHCP Monitor is implemented in software (perhaps open source software), network operators can construct it to operate in a number of ways: it can check against a whitelist and only allow access to known good fingerprints; it can check against a blacklist and only refuse access to known bad fingerprints; or it can keep track of a device's fingerprints over time and revoke access if a device's fingerprint changes. Finally, since the SDN Controller gives operators fine-grained control over the network, operators can decide how to respond to a negative result: they might physically disable the device's port; they could connect the device to an active network profiling tool for further examination; or they might send copies of all packets originating from or destined to a device to a more sophisticated network sensor.

This approach allows for numerous other possibilities in terms of the data used for access-control decisions. For example, an end-point's network access can be linked to its enrollment in an Active Directory environment by having the SDN Controller instruct network devices to provide limited connectivity to new end-points added to the network. Specifically, the network device will only allow a new end-point to communicate with the Active Directory server. This configuration stays in place until Active Directory has made an authorization decision and informed the SDN Controller. The controller can then take the necessary action to allow or deny network connectivity. Organizations utilizing other sources of authorized devices can incorporate them in a similar way.

An important note is that this approach has only a minimal requirement—network devices must implement the OpenFlow protocol. Therefore, it can be supported by many different hardware vendors and many SDN Controllers, and can be deployed much more widely and robustly than today's vendor-specific solutions.

Rapid Response

Two well-known shortcomings of today's networks are that their behavior can be unpredictable and that network changes can cause unpredictable results. In addition, implementing network changes is often a manual process. These factors combine to make troubleshooting network issues a time- and labor-intensive process that cannot keep up with future cyber threats (Feamster & Balakrishnan 2005; Feamster *et al*. 2004). By combining logically centralized control with fine-grained control over network devices, an SDN enables a multitude of troubleshooting capabilities that cannot be implemented in traditional networks. These capabilities enable DCO by allowing operators to more rapidly respond to the changing operational environment. For example, with complex flow-monitoring and flow-modifying tools, a network operator can be alerted when there are complications in the network, where in the network complications have occurred, and what actions need to be taken to resolve the problems. Furthermore, these tools can run periodically to enable continuous, proactive network administration. In the most ideal case, automating the response to the detection of network abnormalities and a periodic check for abnormalities in the network will result in virtually zero downtime for the majority of troubleshooting issues (Sundaresan *et al*. 2010). The following analysis illustrates these benefits by describing a number of example capabilities that could be implemented as applications sitting on top of an SDN Controller.

The first capability allows network operators to determine whether two devices in the network have IP connectivity. In other words, operators can easily figure out whether the network will route packets from one device to the other device, a capability similar to what was proposed by Narayana, Rexford, and Walker (2014). For example, operators might do this by specifying a pair of device identifiers (MAC addresses, IP addresses, host names, etc.) for the SDN Controller to investigate. The SDN Controller then uses its stored topology and device information to determine if the devices are allowed to connect, and it queries all network devices for their currently-programmed traffic-forwarding rules. The SDN Controller analyzes this information to determine if the two devices are allowed to connect and whether the network devices' forwarding behaviors are implementing this policy. If the devices do not have connectivity, the SDN Controller can notify the operator of the specific network devices and forwarding rules that are preventing communication as well as any devices that lack rules to allow communication. Taking this example further, network operators can configure the SDN Controller to continuously evaluate connectivity between mission-critical devices so that any interruption is identified as quickly as possible. This type of monitoring is not practical in today's networks because of the variety of protocols that control network devices' forwarding behaviors, the variety of mechanisms (such as VLANs) network operators use to modify devices' forwarding behaviors, and the vendor-specific nature of both these things. Any practical implementation of this capability requires the centralized, vendor-agnostic, and predictable forwarding behavior that SDN provides.

The second capability builds on the first and allows network operators to automatically enable or disable connectivity in the network by specifying pairs of device identifiers the operators want to affect. Operators might do this in response to network security or failure events. Since the SDN Controller has the network topology information, it can automatically determine what type of forwarding rules to program into network devices to either explicitly allow or disallow connectivity. This capability can also enable DCO by facilitating on-demand transmission of traffic copies to a security team's auditing server in order to log particular events for future analysis and baselining network characteristics.

The last capability allows network operators to determine if the network is configured as expected in a manner similar to the one described by Kazemian, Chang, Zeng, Varghese, McKeown, and Whyte (2013). In particular, the operators can determine if the network is enforcing desired security policies. SDN makes this capability possible because, as described in the previous sections, network operators must explicitly instruct the SDN Controller as to how packets should be forwarded, and the SDN Controller translates these instructions into forwarding rules implemented by the network devices. To validate network behavior, network operators simply reverse this process: they instruct the SDN Controller to extract all forwarding rules from the network devices, and they either use tools to validate that these rules produce the desired behavior or do so manually.

The following two examples illustrate this process. The first example may be labeled as the dynamic-access control scenario. If an operator discovered that a network device did not have a rule to forward DHCP traffic to the monitor, the network is not behaving as desired. In a second example, the removing-network-learning scenario from above, operators might check that forwarding rules for non-ARP packets are only based on destination IP address and conform to the information extracted from the data center's automation tools.

Again, this capability can be configured to run periodically so that network behavior is continuously validated and network operators are alerted as soon as there is something unexpected. As noted above, this type of capability is not practical in today's networks due to the variety of mechanisms affecting network devices' forwarding rules and the vendor-specific nature of those mechanisms.

Open Questions

Just as Software-Defined Networking forces a paradigm shift in traditional networking models, so it also forces change in security models and policies for the organizations that adopt it. Although SDN brings many promising capabilities to DCO, there are many open-ended security questions that organizations will need to address based on their specific networking and security needs. This section highlights some of these questions.

To begin with, organizations implementing SDN will need to decide how granular their rules are for packets inspected on the network and what is given up in order to have deeper levels of security. These are not necessarily new considerations for network security; however, the implementation through SDN frames these issues differently. As an example, the Dynamic-Access Control application detailed above can be written to encompass various security policies. Determining whether a host should connect to a specific port from a specific physical/logical address or some combination is easy to implement in code, and largely a security policy decision. This decision, however, will need to be implemented by the team writing or managing the applications sitting on top of the SDN Controller, and then checked and enforced by auditors. While the process of creating, implementing, and auditing security policies is not a new concept in large organization, SDN provides some extra complexity and flexibility with the granularity of match-based rules.

Next, rule prioritization is also something organizations will need to master. Balancing the operational needs against the security needs of an organization can be difficult. Rules and their prioritization have always been a part of network security. However, given the expansion of capabilities that SDN provides, understanding how the traffic forwarding rules generated by SDN applications affect each other and how they affect the network becomes even more important. In the event that there are multiple SDN applications trying to control forwarding behavior for the same types of network traffic, it becomes very easy to run into issues. One example of this is a firewall application running in parallel with another application that simply forwards traffic. If there is a firewall rule preventing the forwarding of a packet, this same packet may be handled through the other application and compromise the integrity of the network. Perhaps the firewall and forwarding application are merged. Much like today's firewall rules, the ordering of functions and rules within the application can cause unexpected network behavior if misconfigured.

The installation of poorly coded or exploitable SDN applications is also a concern. Security professionals and programmers often fail to communicate properly during the creation of new projects, and it is well documented that the number of bugs in code increases with the complexity and length of code. The importance of secure programming practices becomes that much more important when the software is running the network. Since SDN will be used for DCO, it becomes important to ensure applications function in the ways that they were intended, are put through rigorous testing, and are vetted for security compliance before being added to a network.

Finally, SDN also changes the attack surface of a network. Instead of trying to exploit many individual network devices located throughout the network, attackers now have the SDN Controller as a single point of focus. Communications between the SDN Controller and network devices must be protected with strong access controls that can detect unauthorized connections. The controller must be protected from denial-of-service attacks; consequently, architectures protecting the controller and the creation of applications to monitor the network traffic will be critical for ensuring network security.

Considering these open questions and the new opportunities SDN brings, organizations must ensure their security departments are involved in the migration to SDN so that SDN deployments remain security compliant.

Conclusion

SDN is a fundamental shift in thinking about and architecting networks. Its logically-centralized and fine-grained control enables new defensive capabilities by eliminating network learning, by creating opportunities for dynamic access control, and by facilitating rapid response to changing network conditions. To be sure, network operators can leverage these capabilities while conducting DCO to protect their networks from threats of the future.

References

Caesar, M, Casado, M, Koponen, T, Rexford, J & Shenker, S 2010, 'Dynamic route computation considered harmful', *ACM SIGCOMM Computer Communications Review*, vol. 40, no. 2, pp. 66-71.

Casado, M, Freedman, MJ, Pettit, J, Luo, J, McKeown, N & Shenker, S 2007, 'Ethane: taking control of the enterprise', *ACM SIGGCOM: Proceedings of the 2007 Conference on*

Applications, Technologies, Architectures, and Protocols for Computer Communications Conference, <http://dl.acm.org/citation.cfm?id=1282382>.

Comer, D 2013, *Internetworking with TCP/IP*, vol. 1, Addison,Wesley, Massachussets, United States.

Feamster, N, Balakrishnan, H, Rexford, J, Shaikh, A & Van Der Merwe, J 2004, 'The case for separating routing from routers,' *Proceedings of the ACM SIGCOMM Workshop on Future Directions in Network Architecture*, <http://dl.acm.org/citation.cfm?id=1016709>.

Feamster, N & Balakrishnan, H 2005, 'Correctness properties for internet routing', *Allerton Conference on Communication, Control, and Computing*, <http://www.umiacs.umd.edu/publications/correctness-properties-internet-routing>.

Heller, B, Scott, C, McKeown, N, Shenker, S, Wundsam, A, Zeng, H, Whitlock, S, Jeyakumar, V, Handigol, N, McCauley, J, Zaris, K & Kazemian, P 2013, 'Leveraging SDN layering to systematically troubleshoot networks', *ACM SIGCOMM 2013 Hot Topics in Software Defined Networking*.

Kazemian, P, Chang, M, Zeng, H, Varghese, G, McKeown, N & Whyte, S 2013, 'Real time network policy checking using header space analysis', *10th USENIX Symposium on Networked Systems Design and Implementation*.

Nayak, AK, Reimers, A, Feamster, N & Clark, R 2009, 'Resonance: Dynamic Access Control for enterprise networks', *1st ACM Workshop on Research on Enterprise Networking*.

Naranya S, Rexford J & Walker D 2014, 'Compiling path queries in Software-Defined Networks', *ACM SIGCOMM HotSDN Workshop*.

Open Network Operating System, viewed 18 January 2015, <https://wiki.onosproject.org>.

Open Networking Foundation 2015a, *OpenFlow specification*, viewed 18 January 2015, <https://www.opennetworking.org/sdn-resources/onf-specifications/openflow>.

—— 2015b, *SDN definition*, viewed 18 January 2015, <https://www.opennetworking.org/sdn-resources/sdn-definition>.
OpenDaylight 2015, viewed 18 January 2015, <http://www.opendaylight.org>.

Ryu 2015, viewed 18 January 2015, <http://osrg.github.io/ryu>.

Sundaresan, S, Lumezanu, C, Feamster, N & Francois, P 2010, 'Autonomous traffic engineering with self-configuring topologies', *ACM SIGCOMM Computer Communications Review*.

JOURNAL OF INFORMATION WARFARE

Volume 14 Issue 2 April 2015