VOLUME 7 NUMBER 4
MARCH 2020

CSIAC JOURNAL

A Quarterly Publication of the Cyber Security & Information Systems Information Analysis Center

DATA-CENTRIC ENVIRONMENT

Rise of Internet-Based Modern Warfare "iWar"

Survivability, Availability and Accessibility without Geographic Borders

DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.

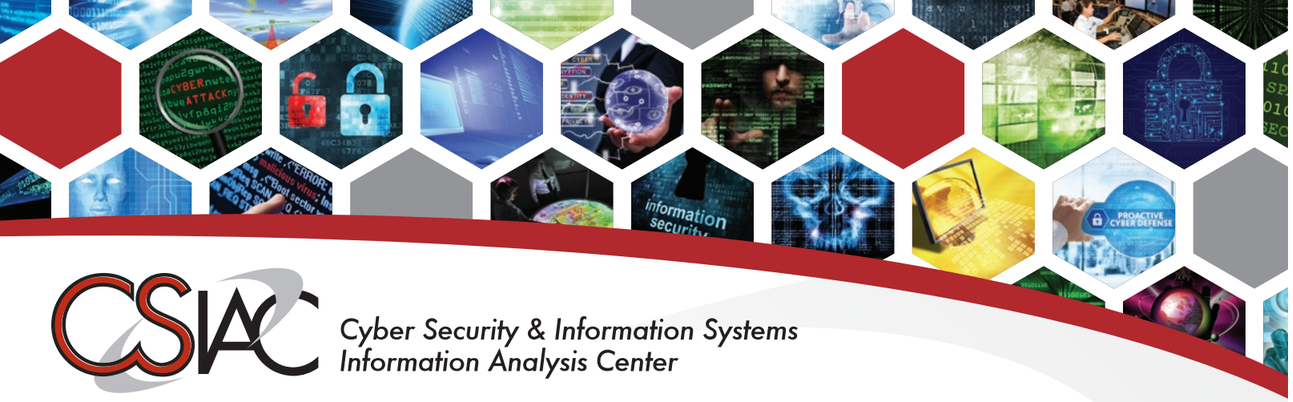

Cyber Security & Information Systems
Information Analysis Center

ABOUT THE CSIAC

As one of three DoD Information Analysis Centers (IACs), sponsored by the Defense Technical Information Center (DTIC), CSIAC is the Center of Excellence in Cyber Security and Information Systems. CSIAC fulfills the Scientific and Technical Information (STI) needs of the Research and Development (R&D) and acquisition communities. This is accomplished by providing access to the vast knowledge repositories of existing STI as well as conducting novel core analysis tasks (CATs) to address current, customer focused technological shortfalls.

OUR MISSION

CSIAC is chartered to leverage the best practices and expertise from government, industry, and academia in order to promote technology domain awareness and solve the most critically challenging scientific and technical (S&T) problems in the following areas:

- Cybersecurity and Information Assurance
- Software Engineering
- Modeling and Simulation
- Knowledge Management/Information Sharing

The primary activities focus on the collection, analysis, synthesis, processing, production and dissemination of Scientific and Technical Information (STI).

OUR VISION

The goal of CSIAC is to facilitate the advancement of technological innovations and developments. This is achieved by conducting gap analyses and proactively performing research efforts to fill the voids in the knowledge bases that are vital to our nation. CSIAC provides access to a wealth of STI along with expert guidance in order to improve our strategic capabilities.

WHAT WE OFFER

We provide expert technical advice and assistance to our user community. CSIAC is a competitively procured, single award contract. The CSIAC contract vehicle has Indefinite Delivery/Indefinite Quantity (ID/IQ) provisions that allow us to rapidly respond to our users' most important needs and requirements.

Custom solutions are delivered by executing user defined and funded CAT projects.

CORE SERVICES

- Technical Inquiries: up to 4 hours free
- Extended Inquiries: up to 2 months
- Search and Summary Inquiries
- STI Searches of DTIC and other repositories
- Workshops and Training Classes
- Subject Matter Expert (SME) Registry and Referrals
- Risk Management Framework (RMF) Assessment & Authorization (A&A) Assistance and Training
- Community of Interest (COI) and Practice Support
- Document Hosting and Blog Spaces
- Agile & Responsive Solutions to emerging trends/threats

PRODUCTS

- State-of-the-Art Reports (SOARs)
- Technical Journals (Quarterly)
- Cybersecurity Digest (Semimonthly)
- RMF A&A Information
- Critical Reviews and Technology Assessments (CR/TAs)
- Analytical Tools and Techniques
- Webinars & Podcasts
- Handbooks and Data Books
- DoD Cybersecurity Policy Chart

CORE ANALYSIS TASKS (CATS)

- Customer tailored R&D efforts performed to solve specific user defined problems
- Funded Studies - $1M ceiling
- Duration - 12 month maximum
- Lead time - on contract within as few as 6-8 weeks

CONTACT INFORMATION

266 Genesee Street
Utica, NY 13502

1 (800) 214-7921

info@csiac.org

/DoD_CSIAC

/CSIAC

/CSIAC

ABOUT THE JOURNAL OF CYBER SECURITY AND INFORMATION SYSTEMS

JOURNAL EDITORIAL BOARD

RODERICK A. NETTLES
Managing Editor
Quanterion Solutions Inc., CSIAC

JOSEPH CAROLI
CSIAC Director
Quanterion Solutions Inc., CSIAC

DR. PAUL B. LOSIEWICZ
Senior Scientific Advisor
Quanterion Solutions Inc., CSIAC

ERIC PATTERSON
Executive Director
SANS Technology Institute

LEAH TREMAGLIO
Student Advisor/Writing Instructor
SANS Technology Institute

DR. KENNETH E. NIDIFFER
Director of Strategic Plans for Government Programs
Carnegie Mellon University
Software Engineering Institute (SEI)

ROBERT J. DEITZ II
Senior Cyber SME in Government / System Integrator Space

GLYN THOMAS GOWING, D.C., PH.D.
Adjunct Professor, Information Assurance, Computer Forensics, Malware
College of Engineering and Computing
Nova Southeastern University
LeTourneau University

MICHAEL R. CONNELLY; CISM
Information Systems Security Group Supervisor/ Chief Strategist Classified Policy & Programs
Johns Hopkins University Applied Physics Laboratory

SHELLEY STOTTLAR
Graphic Designer
Quanterion Solutions Inc., CSIAC

ABOUT THIS PUBLICATION

The Journal of Cyber Security and Information Systems is published quarterly by the Cyber Security and Information Systems Information Analysis Center (CSIAC). The CSIAC is a Department of Defense (DoD) Information Analysis Center (IAC) sponsored by the Defense Technical Information Center (DTIC) and operated by Quanterion Solutions Incorporated in Utica, NY.

Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the CSIAC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the CSIAC, and shall not be used for advertising or product endorsement purposes.

ARTICLE REPRODUCTION

Images and information presented in these articles may be reproduced as long as the following message is noted:

"This article was originally published in the CSIAC Journal of Cyber Security and Information Systems Vol.7, No 4"

In addition to this print message, we ask that you notify CSIAC regarding any document that references any article appearing in the CSIAC Journal.

Requests for copies of the referenced journal may be submitted to the following address:

Cyber Security and Information Systems
266 Genesee Street
Utica, NY 13502

Phone: 800-214-7921
Fax: 315-732-3261
E-mail: info@csiac.org

An archive of past newsletters is available at **https://www.csiac.org/journal/.**

*To unsubscribe from CSIAC Journal Mailings please email us at **info@csiac.org** and request that your address be removed from our distribution mailing database.*

COVER PHOTO

Cover Graphic Composite: Shelley Stottlar, Quanterion Solutions Inc., **Featuring Deposit Photos Stock Images:** by monsit, pyty, SergeyNivens, and macrovector.

JOURNAL OF CYBER SECURITY AND INFORMATION SYSTEMS

Data-Centric Environment - Rise of Internet-Based Modern Warfare "iWar"
Survivability, Availability and Accessibility without Geographic Borders

Photo Graphic Composite: Shelley Stottlar, Quanterion Solutions Inc., **Featuring Deposit Photos Stock Images:** by sdecoret, everythingposs, IgorVetushko, alexmillos, and GoodLuckWithUs.

EVALUATION OF COMPREHENSIVE TAXONOMIES FOR INFORMATION TECHNOLOGY THREATS

By: **Steven M. Launius**, SANS Technology Institute, Candidate, MS Information Security Engineering

CATEGORIZATION OF ALL INFORMATION TECHNOLOGY THREATS CAN IMPROVE COMMUNICATION OF RISK FOR AN ORGANIZATION'S DECISION-MAKERS WHO MUST DETERMINE THE INVESTMENT STRATEGY OF SECURITY CONTROLS.

While there are several comprehensive taxonomies for grouping threats, there is an opportunity to establish the foundational terminology and perspective for communicating threats across the organization. This is important because confusion about information technology threats pose a direct risk to an organization's operational longevity. In order for leadership to allocate security resources to counteract prevalent threats in a timely manner, they must understand those threats quickly. A study that investigates categorization techniques of information technology threats to non-technical decision-makers through a qualitative review of grouping methods for published threat taxonomies could remedy the situation.

INTRODUCTION

A modern organization's operations depend largely on information technology (IT). Ubiquitous adoption of IT due to technological advancements creates both efficiencies and vulnerabilities in an organization's operations. Physical threats to IT infrastructure from both human and environmental sources have remained mostly consistent over time. The continuous development of IT systems for exchanging, processing, and storing information introduces many weaknesses. Criminals, activists, nation-stations, and other adversaries are increasingly successful at attacking these systems to accomplish their objectives. Many organizations are adopting Cyber Threat Intelligence (CTI) to address the increase in adversarial cyber threats. Since the primary use of CTI is the sharing of an adversary's activities, several taxonomies and ontologies exist for maintaining a common lexicon within and between organizations.

However, in addition to nefarious humans, sources of IT threats may also be accidental, environmental, political, or economic. Leadership must evaluate risk to IT by assessing the likelihood of threat events from all of these sources and their impact on the organization. Risk management professionals from the information security community have published comprehensive taxonomies for grouping threats events. Each taxonomy presents a hierarchy of discrete threat event groups with succeeding levels providing terms with more detail. Categorization and definitions of terms for threat events support communication with decision makers who must select a course of action to counter a threat.

A threat taxonomy can improve communication in two ways. First, language barriers between professionals with different expertise can be broken down into clear definitions for IT threats. As mass media quickly spreads news of IT failures, like cyberattacks or data breaches, a foundation of terms can help

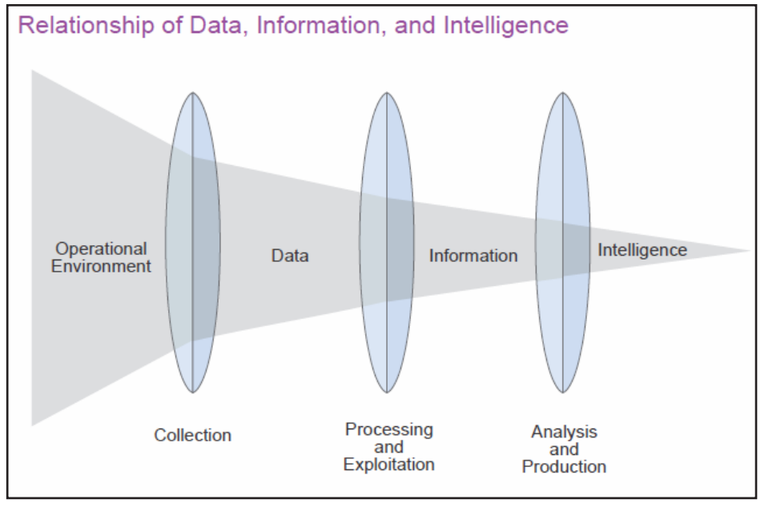

Figure 1: *Relationship of threat data, information, and intelligence.*

decision-makers understand the active threats. Second, an ordered taxonomy structure of the entire IT threat landscape enables analysis and assessment at various granularities. Comparing the risk of high-level threat categories can empower leadership to make the right decisions to protect their organization.

COMMUNICATING THREAT

Threat Language

Language is an intricate cognitive process requiring an agreement of standard definitions for effective communication. While the English language has broadly held standards, there are many deviations that can present communication problems. In particular, slang differences occur at many levels:

- National: Americans live in *apartments*, while Brits live in *flats*.
- Regional: Soda, pop, coke, and soft drink are all terms for a sweetened carbonated beverage.
- Local: In Texas, a nag is called a *worrit*.
- Professional: In the health profession, a virus is a microorganism that infects living cells to live and reproduce itself and causes human illness (Definition of Virus, 2018). In the IT profession, a virus is a hidden, self-replicating section of computer software, usually malicious logic, propagating by infection of another program (Glossary of Security Terms, 2018).

Adhering to standard definitions for threat terms can improve comprehension of the dialog between echelons in any organization. There is no authoritative source for IT threat terms, but there are several glossaries or lexicons of security terms published by a variety of governing bodies. The United States (U.S.) government alone has many sources including:

- Department of Defense (DoD) - Dictionary of Military and Associated Terms,
- Department of Homeland Security (DHS) - Risk Lexicon,
- National Institute of Standards and Technology (NIST) - Glossary of Key Information Security Terms,
- Committee on National Security Systems (CNSS) - Glossary, and
- National Initiative for Cybersecurity Careers and Studies (NICCS) - A Glossary of Common Cybersecurity Terminology.

Many information security organizations also maintain security term definitions:

- SysAdmin, Audit, Network, and Security (SANS) Institute - Glossary of Security Terms,
- Information Systems Audit and Control Association (ISACA) - Cybersecurity Fundamentals Glossary,
- International Organization for Standardization (ISO) - Search for Terms & Definitions,
- Internet Engineering Task Force (IETF) Trust - Request for Comments (RFC) 4949 Internet Security Glossary,
- Information Technology Infrastructure Library (ITIL) v3 - Foundation Course Glossary.

There is some agreement between definitions, but it is not reasonable for non-technical professionals to learn the abundant terms and nuances of each. A smaller set of organizational-wide IT threat terms are necessary for more business-oriented professionals.

A discrete set of IT threat categories with standard definitions can increase communication and support risk reduction. Information security operations provide analysts with a rich vocabulary of cyber threat terms and a structure for appropriately characterizing attacks. CTI and incident response operations describe and analyze an attack in great detail to support threat hunting, sharing, and governance of information security operations. A taxonomy of IT threat terms can provide appropriate categories at various levels of granularity to aid threat analysis, risk assessments, and ultimately decision-making. Capturing and organizing unstructured threat information through CTI and incident response activities requires a standard set of threat terminology. Reports and metrics with a common set of terms can speed comprehension of the threats and incident response times. Business unit management and organizational leadership can more quickly understand the greatest threats to their organization after reviewing threat reports and metrics with standard terminology.

Since organizational leadership makes decisions based on risk, threat terms must be able to support risk management. All businesses must balance risk with reward, but severe consequences may result from misunderstanding the risk. An accurate depiction of the threats to information technology is vital for leadership to make appropriate decisions. Organizations in many industries use a variety of risk frameworks that may be threat-, vulnerability-, or asset-based. Regardless of the risk framework type, the quantities of threats should be commensurate with the maturity of the organization's risk management. Listing every possible hazard in an immature implementation of a risk framework can overwhelm risk analysis and bring the process to a halt. The risk management process should use threat categories appropriate for the maturity of the organization's risk assessment.

Threat Taxonomy for Cyber Threat Intelligence

CTI was born from the application of military intelligence doctrine to data analysis of cyberattacks. The DoD describes the intelligence process as a cycle of phases: direction, collection, processing, analysis, dissemination, and feedback (JP 2-0, 2013). While represented as a cycle, the steps may happen concurrently or may be skipped entirely depending on the situation. The intelligence cycle prescribes the process for collecting threat data and transforming it into threat intelligence. Brian P. Kime's article, "Intelligence Preparation of the Cyber Operational Environment" relates the DoD intelligence cycle to information security by presenting a collection method for threat data from IT infrastructure (Kime, 2016). Figure 1 shows the transformation of threat data into information, via structure and context, then into intelligence, via analysis, as it flows through the intelligence cycle phases.

Structuring data to produce information is precisely where an IT threat taxonomy fits into CTI. A threat taxonomy sits on top of the available standards and ontologies for capturing threat data.

There are several CTI standards for modeling, storing and sharing threat data from cyberattack investigations. These standards capture indicators of compromise (IOC) or attacker tactics, techniques, and procedures (TTP). IOC are the easy-to-modify artifacts with the context pertinent to a cyberattack, such as file hashes of malicious program files or domain names of phishing websites. TTP describe the actions, skills, methods, or modus operandi (MO) adversaries use to accomplish their goals. Threat models help relate IOC and TTP to each other for an illustration of the overall attack process and objectives during analysis. Robert M. Lee and Mike Cloppert describe threat modeling, such as Cyber Kill Chain and Diamond models, as an intrusion analysis technique for understanding threats and prioritizing defensive efforts that drive security (Lee, 2016). Organization and collection of the similar actions and techniques of cyberattacks facilitate sharing between industry partners and government bodies. Greg Farnham's paper on "Tools and Standards for Cyber Threat Intelligence Projects" (Farnham, 2013) presents and defines many CTI standards for an evaluation of a project management process. Those relevant for storing and sharing TTP include Structured Threat Information eXpression (STIX), Open Indicators of Compromise (OpenIOC) framework, and Collective Intelligence Framework (CIF).

"A discrete set of IT threat categories with standard definitions can increase communication and support risk reduction."

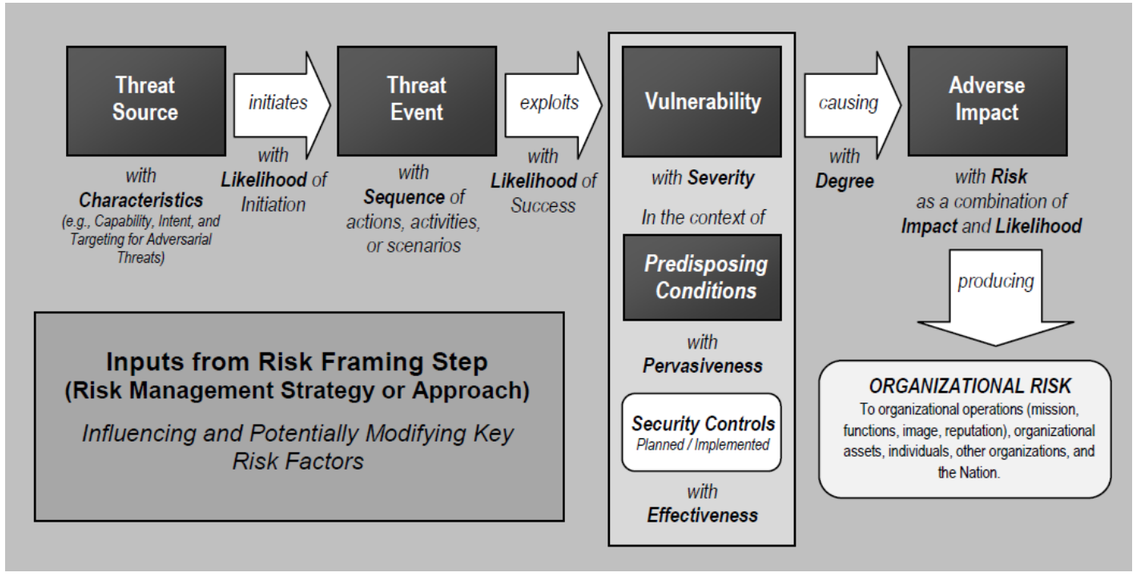

Figure 2: *NIST 800-30 generic risk model with key risk factors.*

While CTI standards provide structure for comprehensive threat analysis by subject matter experts, they often lack general groupings necessary for decision-makers to understand threats. According to the SANS 2017 CTI Survey (Shackleford, 2017), CTI standards have seen widespread adoption within CTI programs since Farnham's article was published. STIX consists of even more granular CTI standards. The Common Attack Pattern Enumeration and Classification (CAPEC) is a standard for describing cyberattack patterns (MITRE, 2017) that fits into STIX. CAPEC has 508 terms to portray all possible attack patterns. STIX and CAPEC are examples of the intricate threat detail capable with CTI standards. These capabilities aid threat analysis, but a higher-level perspective supports strategic CTI products.

CTI has three levels of analysis with a different purpose and audience for each: strategic, operational, and tactical. The operational and tactical levels of intelligence analysis concentrate on tracking and sharing attacker IOC and TTP with the CTI standards as previously explained. Analysis at the strategic level of CTI requires the same threat information, but addresses the overall risk to the organization by answering questions about cyber threats from leadership. The "Operational Level of Cyber Intelligence" published in the International Journal of Intelligence and CounterIntelligence provides an overview of these levels suitable for this discussion (Mattern, 2014). Strategic level intelligence "… pertains to an organization's general direction, specific goals, and resource allocation in service of its mission, as guided by the highest-level executive or command entity." Strategic intelligence analysis includes comparing security resources to trend changes in threats over time. At this level, intelligence analysis informs business units about the most likely threats to impact operations and the resources necessary to reduce this risk. A threat taxonomy supports strategic intelligence analysis with a consistent threat perspective to satisfy the needs of organizational leadership.

Within the private sector, CTI operations concentrate on operational and tactical levels of analysis. The SANS Institute sponsors an annual survey of CTI since 2015 that demonstrates a focus on operational and tactical intelligence analysis, specifically on IOC. Comparison of the last three reports reveals a growing adoption of CTI with security tools primarily designed for identification, collection, or correlation of IOC. According to the 2015 survey, CTI improves security and response by increasing visibility into attack methodologies, cited by 63% of respondents, and by increasing incident response times, cited by 51% of respondents (Shackleford, 2015). The top three use cases in the 2016 survey were blocking malicious IP addresses or domain names at the firewall, adding context to incidents, and identifying malicious activity through DNS logs (Shackleford, 2016). The 2017 survey indicates that most organizations have dedicated CTI teams for collecting and processing CTI data (Shackleford, 2017).

These same studies also show the lack of application to strategic analysis. In the 2016 survey, more than half of the respondents said CTI is important to risk prioritization and decision making, but

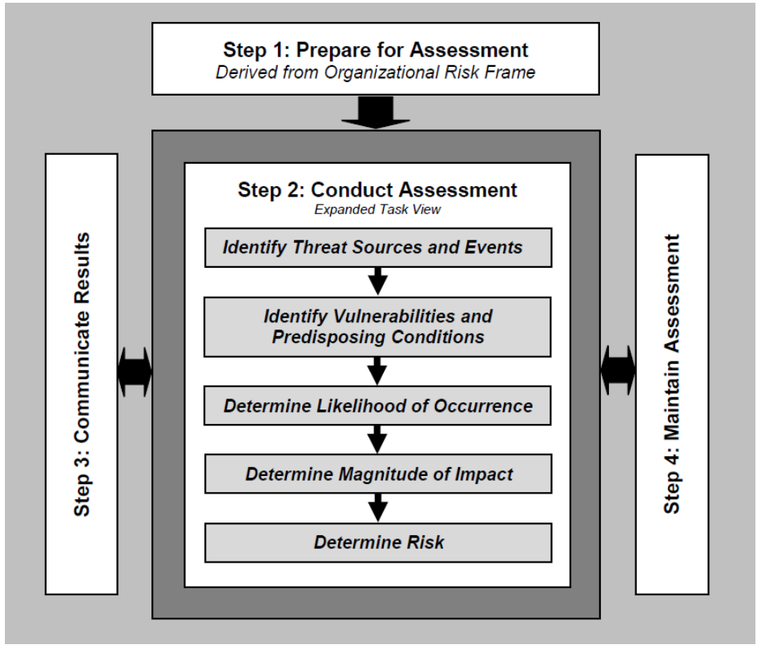

Figure 3: *Risk assessment steps from NIST 800-30*

the 2017 survey lists "budget and spending prioritization and decisions" lowest among the use cases for CTI. Therefore, it is not surprising to see the primary skills for strategic analysis reporting, writing, presentation, and oral communications at the bottom of the skills list for CTI analysts in the 2017 survey. The survey respondents indicate the value of CTI is from an increase in preventing attacks and responding to attacks. However, CTI does not appear to be affecting strategic-level decisions. An inability to communicate with business terms the sources threating specific business operations and the appropriate security measures to reduce this risk are the likely reasons why CTI is not influencing leadership.

Standard threat categories and terms in a taxonomy of all IT threats can assist analysis for producing strategic-level intelligence. Many publicly available intelligence sources produce unstructured reports. These intelligence sources frequently describe the same threat with various synonyms or attack terms. There is little agreement between sources of the names given to adversaries, malware, or attack techniques. Aggregation of the threat components, while consuming intelligence from a variety of sources, supports automated analysis methods. A threat taxonomy can help match these external reports to internal incidents. Organizations can predict future adversary actions by identifying attack patterns when threat modeling has a standard terminology. Revealing trends in attack vectors and adversary methods is possible when analyzing cyberattacks with a threat taxonomy. This type of analysis is useful for risk management because identifying the most likely threats helps prioritize remediation. Threat frameworks with detailed ontologies of threat information are difficult to use in risk analysis. Given the number of possible actors, actions, targets, and consequences for every threat, the list of possible threat events may total in the thousands or more. Governance, risk, and compliance (GRC) tools can provide an organization with automation of risk assessment calculations for complex threats. However, GRC tools are not available within every organization or may not support CTI standards. In the absence of these tools, scripts or macro-enabled productivity software can provide sufficient automation of workflow to produce CTI products usable in a risk assessment. Grouping threat information into a taxonomy provides a finite set of threat scenarios, so the risk analysis process does not overwhelm available resources.

Threat Taxonomy for Risk Assessments

The rich threat information in CTI can support information security risk frameworks, but assessing non-adversarial threats is also important. An adversarial threat taxonomy in a CTI program needs to be merged with non-adversarial threats, like environmental or human mistakes, in a risk assessment to communicate the level of risk across all threats facing an organization's information services. Risk frameworks from organizations like NIST, ISO, US-CERT, ISACA, and others use likelihood estimates for both adversarial and non-adversarial threats in the assessment process. The Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) is a risk management methodology from Carnegie Mellon University and US-CERT. OCTAVE Allegro (the most recent version) is an information asset-based assessment methodology which uses simple qualitative assessments of threat profiles. ISACA's latest version of Control Objectives for Information and Related Technologies (COBIT) 5 for Risk addresses risk of enterprise IT governance in the form of principals and guidance. To demonstrate the integration of a threat taxonomy into a risk framework the NIST's Risk Management Framework (RMF) provides a useful open and mature framework (NIST SP 800-37, 2010). NIST's Guide for Conducting Risk Assessment (NIST SP 800-30, 2012) provides important concepts and processes for implementing the RMF and describes where a threat taxonomy interacts with the risk assessment. Identifying, estimating,

and prioritizing information security risks are the function of a risk assessment.

Threats are one common risk factor NIST's risk assessment methodology identifies for assessing and relating risks in a model. The risk factors define the characteristics for determining risk levels that are essential for communicating problematic situations. Definitions for risk factors are informed by an organization's risk management strategy or during risk framing if a strategy does not exist. The other key risk factors seen in Figure 2 include vulnerability, impact, likelihood, and predisposing condition. Threats break down into threat sources that cause threat events. A threat event has the potential to negatively impact an organization's operations or assets through the loss of confidentiality, integrity, or availability of information or information systems. A threat source is the "intent and method targeted at the intentional exploitation of a vulnerability or a situation and method that may accidentally exploit a vulnerability" (NIST SP 800-30, 2012). NIST's comprehensive overview of threat sources includes:

- Cyber or physical attacks
- Human errors
- Failure of resources
- Environmental disasters, accidents, or failures

NIST prescribes a four-step risk assessment process, illustrated in Figure 3, for preparing, conducting, communicating results, and maintaining a risk assessment. Organizations define and use the threat taxonomy in the first two steps of the risk assessment process. During Communicate Results in the third step, the report and metric products sent to leadership should use this same threat terminology.

In preparation for the risk assessment, organizations can define a threat taxonomy in the first step as part of risk fra ming. Identifying the main assumptions relevant to risk assessments is one of the tasks which enables the RMF to clarify risk models and increase repeatability of results. Two of the key assumption areas are threat sources and events. The level of detail chosen for threat sources and events will establish the set of possible threats available when identifying the relevant threats to the organization in the Conduct Assessment step.

Another crucial assumption area for risk assessments is the analytic approach for characterizing threat sources and events. The analytic approach consists of both the assessment type (i.e. quantitative, qualitative) and analysis type (i.e. threat-, asset-, of vulnerability-orientated). A many-to-many relationship exists among threat events and sources, therefore levels with greater detail increases the complexity of the risk assessment. A threat taxonomy categorizing all possible threat sources and events with varying levels of granularity can allow an organization to move from less to more detail as their risk management program matures.

"The taxonomy will only categorize the threat event component, but events must include activities from both human and environmental threat sources."

COMPREHENSIVE THREAT TAXONOMIES

A taxonomy is an ordered classification system, often hierarchical, where each parent tier is a grouping of the terms characterizing its child tier. The terms each taxonomy uses for the hierarchical levels are slightly different but serve a similar purpose. Descriptive terms for the top-level of a taxonomy may include class, top-tier, or high-level. Terms for the second level of a taxonomy may include family, threats, or subclasses. The designations for taxonomies with a third level consist of elements or threat details. The terms and structure of each taxonomy used in this research can be found in Appendix A.

Several institutions have created comprehensive threat taxonomies for IT systems. A comprehensive threat taxonomy will have several features. A simple hierarchal structure is necessary where the top tier has no more than ten categories. This discrete set of categories must work to organize events, activities, situations, or contexts from diverse sources of threats encompassing both adversarial and non-adversarial threats. The taxonomy will only categorize the threat event component, but events must include activities from both human and environmental threat sources. The subcategories should include more detail than the higher-level groups with definitions for the terms. Definitions of all threat categories are valuable for creating consensus among the professionals who will work with the taxonomy.

Most of the qualifying taxonomies are incomplete as work on them has only begun within the last few years. Each taxonomy has a different goal and purpose that shapes the categories selected for it. For example, the business operational threat categories of Carnegie Mellon University's taxonomy use business-orientated terms including *people*, *process*, *technology*, and *external*. Mapping these taxonomies should be straightforward with any of the published security control recommendations, like NIST 800-171. The threat taxonomies are primarily for organizations with threat intelligence capabilities to provide probability estimates for threat activities during risk management. In addition to a review of the goal and purpose of the taxonomies, a short analysis of their qualities will reveal their strengths and weaknesses.

Open Threat Taxonomy

The goal of the Open Threat Taxonomy (OTT) was to create a shared and comprehensive set of information system threats that organizations may face. James and Kelli Tarala, authors of the OTT and owners of the security firm Enclave Security, released version 1.1 as an open source tool in October 2015. The OTT defines a threat as "… the potential for a threat agent to cause loss or damage to an information system" (Tarala, 2015). Part of the complexity of defining threats comes from the components that compromise a threat. The OTT lists these components as threat source or agent, threat action, threat target, and threat consequence. Tarala describes the relationship of these components as, "A threat source will most often perform a threat action against a threat target, which leads to threat consequences" (Tarala, 2015). This taxonomy only describes threat actions, but uniquely includes a priority ranking for each action. A one to five scale ranks the priority of each threat, where priority should go to threats with a higher rank. Threat models and attack observations from contributors to the OTT help establish the priority scores and "should be viewed as consensus guidance" (Tarala, 2015).

The OTT covers most of the pertinent threats to information system operations without forgetting most of the non-technical dangers. The OTT categorizes threats by their nature and by the extent to which they impact the confidentially, integrity or availability of information systems. This taxonomy has a total of 75 threat actions broken down into four main categories:

- Physical Threats
- Resource Threats
- Personnel Threats
- Technical Threats

Definitions for each category elaborate on the nature of each threat group. However, the threat actions do not have definitions, only clear descriptive terms. Even though there are short action phrases, an audience's experience could lead to ambiguous interpretations of the terms. The small set of threat categories describes actions that can cause damage to information systems. Adverse impact is defined as threats to confidentiality, integrity, or availability of each category. Therefore, many of the threat actions have an adversarial perspective. This grouping perspective results in a concentration of threat actions within the Technical Threats category as technical vulnerabilities in information systems are numerous. The categorization of all possible threat sources is incomplete, as capturing legal threats does not appear to be possible in the OTT.

The holistic coverage of information systems threats from OTT can provide broad risk comparison across an organization. The OTT works well with risk frameworks that consider inherent and residual risks separately. This is due to priority ranking scores a group of industry experts assigns to each OTT threat action. This ranking system allows an organization to prioritize one threat over another when it must choose between investing in resources to mitigate threats with the same likelihood of occurring. Besides the threat actions, the taxonomy does not address other threat components or help with identifying mitigation controls. Mapping the threat actions to specific security controls, such as NIST 800-53, could assist in completing a risk assessment.

"The holistic coverage of information systems threats from OTT can provide broad risk comparison across an organization."

ENISA Threat Taxonomy

In January 2016, the European Union Agency for Network and Information Security (ENISA) published a taxonomy as an aid for threat information collection and consolidation (ENISA, 2016). The ENISA Threat Taxonomy (ETT) defines Cyber Threats as "… threats applying to assets related to information and communication technology." ENISA's purpose for its taxonomy is to provide definitions for threat terms with a possibility of rearranging its structure. The ETT was designed as an analysis mechanism for collecting and sorting threat information.

The ETT provides a unique view of possible threat actions, but without the consistency and clarity found in other taxonomies. The eight or nine, depending on the version, high-level categories of the ETT are a mixture of consequences and intentions for the 75 total threats actions. The *high-level threats* include:

- Physical Attack
- Unintentional Damages
- Disasters
- Failures / Malfunction
- Outages
- Eavesdropping / Interception / Hijacking
- Nefarious Activity / Abuse
- Legal

The *threats* and *threat details* make up the next two levels of the ETT creating one of the most detailed threat taxonomies. While there is an expectation of change for different versions of a taxonomy, the lack of consistent relationships and accurate definitions throughout the ETT detract from the purpose of a taxonomy. One inconsistency is the alternate terms for three of the high-level threats. The ETT uses a slash symbol to expand the terms of these categories instead of using a single term and definition like the other categories. The high-level threat definitions do not support mutually exclusive categories. For example, the Eavesdropping threat has a definition that fits into the Nefarious Activity threat, but these categories exist at the same level. Additionally, several of the threats and threat details include the threat source or intentions in the description restricting its scope, which will lead to necessary revisions in the future. The

lack of delineation between threat events and sources also causes ambiguous classification of a threat into multiple categories. Such a classification supports complex relationships in threat ontologies, but conflicts with the simplifying purpose of a taxonomy. Similar to OTT, the ETT adversarial threats focus on attacker actions that can negatively impact information systems but disperses them into more high-level threat categories. The ETT brings legal threats clearly into consideration with the inclusion of a Legal category for regulations, changes in law, and the political environment.

NIST Risk Assessment Threat Exemplary

The appendix within NIST's Guide for Conducting a Risk Assessment includes exemplary threat events that provide a sample threat taxonomy. NIST's risk model decomposes threats into a source and event for analysis of a single threat. A series of threat events can create a threat scenario that NIST defines as "a set of discrete threat events, attributed to a specific threat source or multiple threat sources, ordered in time, that result in adverse effects" (NIST SP 800-30, 2012). Multiple events from the same threat source or multiple threat sources executing the same threat event may compromise threat scenarios. These scenarios can result in many granular circumstances; therefore, a mature risk management process is necessary to handle the numerous scenarios that result from this analysis. An organization need only to assess the relevant threat events when there is an adversary with intent or capability to initiate an attack.

For consistent comparisons with other taxonomies, the evaluation will only include the NIST exemplary threat events. The NIST model breaks all threat events into two high-level categories:

- Adversarial
- Non-adversarial

The two-level hierarchy in this taxonomy results in a concentration of threat events for the adversarial category. The second-level categorizations of adversarial threat events are similar to the stages in the Lockheed Martin kill chain model (Lockheed) that characterize adversarial TTP. These stages of a cyberattack include reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objectives. The NIST guide references the MITRE Corporation's CAPEC for characterizing cyberattacks with greater detail (CAPEC, 2017). These adversarial attack patterns describe possible methods for exploiting information systems from an attacker's perspective. The adversarial events categorized by the kill chain stages can be useful for mapping with security controls, like NIST SP 800-53. There are far fewer non-adversarial threat events in NIST's taxonomy and, therefore, no additional subcategories for this type of threat. The non-adversarial category is also lacking many of actions found in other taxonomies for unintentional, accidental, legal, or other non-malicious actions. This sample threat taxonomy may not be useful for an organization unless the threat categories are extended.

"The operational risk terms from Risk Lexicon from DHS (DHS, 2008) are the basis for the threat categories."

Taxonomy of Operational Cyber Security Risks

A comprehensive threat taxonomy from Carnegie Mellon University is one of the oldest available. In 2010, the Software Engineering Institute (SEI), a federally funded research and development center based at Carnegie Mellon, produced the first version of the Taxonomy of Operational Cyber Security Risks (TOCSR) (CMU/SEI, 2014). The taxonomy was updated in 2014 to map with the security and privacy controls in Version 4 of NIST SP 800-53. This taxonomy categorizes instances of *operational cyber security risks* defined as "operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems." The purpose of TOCSR is to provide a tool for identifying all the operational cyber security risks within an organization.

The concise terms and categorization method of TOCSR produces a taxonomy that can assist in risk assessment activities. The primary emphasis of the categorization method is on operational risks to information systems. The TOCSR characterizes threats from a business risk perspective, instead of a threat source perspective as in the other threat taxonomies. This results in categories of threats actions for people, process, and technology. This method results in four top-level categories that SEI calls classes:

- Actions of people
- Systems and technology failures
- Failed internal processes
- External events

In SEI's terminology, each class decomposes further into subclasses and elements.

The operational risk terms from Risk Lexicon from DHS (DHS, 2008) are the basis for the threat categories. While this taxonomy aligns with SEI's OCTAVE method for risk assessments, threat taxonomies are not exclusive to one risk framework. Representation of a complete attack scenario may require a combination of TOCSR threat categories. For practical implementation in the NIST risk assessment, threat elements from multiple classes or subclasses will compose a single scenario. For example, a *software* flaw present in a production web application due to inadequate *testing* could be a result of any element under *actions of people*. SEI provides a mapping to the security guidelines in NIST 800-53.

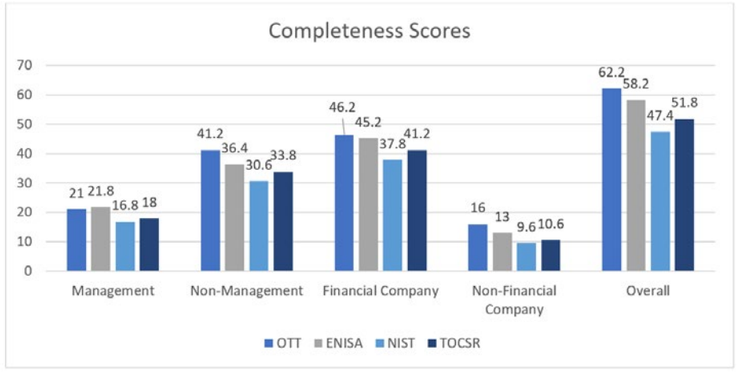

Figure 4: *Completeness scores of each threat taxonomy by respondent groups.*

Other Threat Taxonomies

There are several other published taxonomies for adversarial threats or intelligence sharing. As the need for a taxonomy arose with the formal gathering and sharing of cyberattack information, the work of developing suitable taxonomies is still ongoing. Many organizations only address the most prevalent threats or create taxonomies for specific threats. In either case, these taxonomies are not suitable for an organization-wide taxonomy of threats.

There are many more adversarial-centric threat taxonomies which provide a multitude of options for categorizing the variety of malicious human cyber activities. However, these do not allow comparisons with environmental threats and therefore do not meet the criteria for consideration of a comprehensive threat taxonomy. The aforementioned CAPEC is one such taxonomy of cyberattack patterns by MITRE. Another adversary-centric taxonomy comes from the US government called the Cyber Threat Framework (CTF). The CTF was designed to improve communication between cyber experts and senior leadership across many departments throughout the intelligence community (ODNI, 2017). The variety of threat models in use at different government agencies made sharing cyber threats difficult because of different terminology that was highly technical. Many other CTI standards can map into the four stages of adversary cyberattacks in the CTF. The Office of Director of National Intelligence provides a lexicon for the CTF that equates to a threat taxonomy. The flexible design of the framework allows different views of same adversarial threat information for diverse audiences. One final example of an adversarial threat taxonomy comes from Agari, a secure email exchange company, specifically for cyberattacks against messaging systems (Jakobsson, 2017). The taxonomy breaks down the steps for attacking an email system that was extended to all types of messaging systems, including instant messaging. The scope of these adversarial threat taxonomies is too narrow for organizing a comprehensive set of threats meant for an organization-wide risk assessment.

Researchers at Georgetown University are creating a taxonomy for the existing threat intelligence sharing standards. This cyber threat intelligence information sharing exchange ecosystem (CyberISE) (Burger, 2014) is a classification system for CTI sharing standards. Eric Burger's research presents the structure and relationship to other information sharing technology. The organization of the CyberISE has five top-level categories in a layered model, mimicking the Open Systems Interconnection (OSI) model. The two lower layers address the exchange and authorization of information sharing, while the three upper layers categorize the information exchange. The *Indicators* layer holds the details of an incident or cyberattack. The *Intelligence* layer contains actions to perform when detecting indicators or assessing threats. The 5W's layer comprises the types of questions to ask incident indicators to determine whether an attack is occurring. Since the CyberISE model is for characterizing the existing information sharing standards, it is not an appropriate taxonomy for the categorization of threat information.

The Cambridge Risk Framework is a global threat taxonomy for business operations by the University of Cambridge. The report A Taxonomy of Threats for Complex Risk Management (Coburn, 2014) presents the Cambridge Taxonomy as a taxonomy of macro-catastrophe threats. The basis for threat categorization is extreme events with potential to cause damage or disrupt global social and economic systems. Extreme events have a large impact on global trade and commerce across multiple continents.

Cambridge's development methodology includes a review of historical events and disaster catalogs to create a hierarchy structure of 5 primary classes, 11 families, and 55 types. The report includes definitions for the five classes: Finance & Trade, Geopolitics & Society, Natural Catastrophe & Climate, Technology & Space, and Health & Humanity along with their corresponding families. Insurance risk management is a primary application of the Cambridge Taxonomy. Secondary functions involve risk management of business operations, national security, and finances. While extreme events will have some impact even to small business operations, the likelihood of a global macro-catastrophe event occurring should be overshadowed by more likely, local catastrophes for most businesses. Additionally, the other selected comprehensive threat taxonomies are

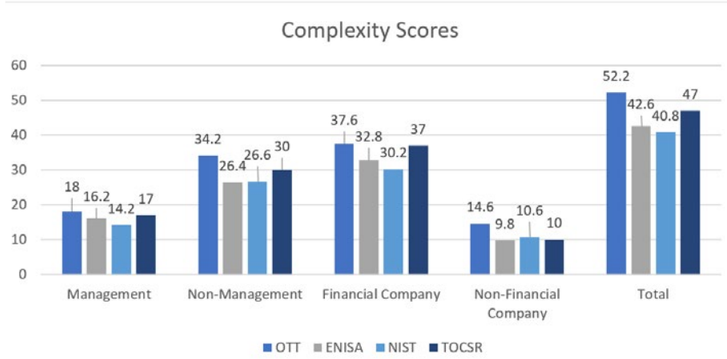

Figure 5: *Complexity scores of each threat taxonomy by respondent groups.*

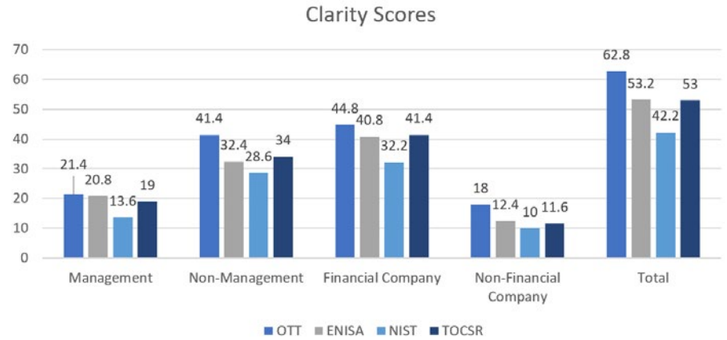

Figure 6: *Clarity scores of each threat taxonomy by respondent groups.*

IT-centric to the effects of threat events. Therefore, the Cambridge Taxonomy was not included in this research evaluation, but global organizations may want to consider it. Organizations of any size may choose to consider this threat taxonomy by redefining catastrophes and extreme events to include disasters at any scale.

THREAT TAXONOMY EVALUATION

In today's environment, cyber defenders are challenged with exploring their threat intelligence capabilities and understand their position against the ever-changing cyber threat landscape. This research evaluation of threat taxonomies uses a qualitative research survey. A qualitative research methodology best supports results dependent upon personal opinions and diverse perspectives. The primary survey focuses on a large financial services company. The risk management department of this company agreed to receive the survey. Responses from this source were plentiful with a total of 61 respondents, labeled as 'Financial Company' in the analysis. An attempt was made to obtain diverse perspectives outside of the Financial Services industry by posting the survey to several social networking forums including information security and educational email list serves as well as professional networking websites. Unfortunately, the response from these sources was much smaller with a total of 23 respondents, labeled as 'Non-Financial Company' in the analysis. The survey began by asking all respondents their industry and job role. To represent different perspectives the analysis compares responses from four groups: Management, Non-Management, Financial Company, and Non-Financial Company. Presentation of the terms and structure of each taxonomy were straightforward, but minor changes were necessary due to formatting restrictions in the survey tool.

There is a potential for respondents to favor the presentation format of a taxonomy while presenting the survey. Authors of the taxonomies use various formatting styles in publications, but to avoid any bias the survey has a consistent table formatting for all the taxonomies. Presentation of the taxonomies took the form of uniform tables. The top-tier categories are set in header rows with the same blue color background. The second tier follows in the next row with categories in a bold font and specific threat actions in a bulleted list for the third tier. The survey mitigates further bias by presenting the taxonomies in a randomly chosen order.

The survey includes only the first two levels of the more complex taxonomies to keep respondent review time to a minimum. Both NIST and ENISA have three or more tiers that can be both overwhelming and tedious to review. The top two tiers list all the major threat categories for each taxonomy. However, the taxonomies presented without the third tier are likely to have lower ratings for completeness. This effect can be even more profound when the clarity of the top tier categories is low, indicating a respondent would not be able to infer the types of threats in a category without them explicitly listed. Reducing the threat actions in the OTT was also necessary for repetitive actions using similar methods. For example, reducing the eleven Application Exploitation actions with different attack methods into a single threat action in the Technical Threat category saves review time without detracting from the threat event. The length of the taxonomies was a likely factor in completing the survey. Fifteen percent of the respondents failed to complete review of all four taxonomies. The OTT had the most responses with about ten more than the other taxonomies. See Appendix B for a complete view of each taxonomy in the same presentation format and order.

The characteristics chosen for evaluation include completeness, complexity, and clarity. These traits were chosen for

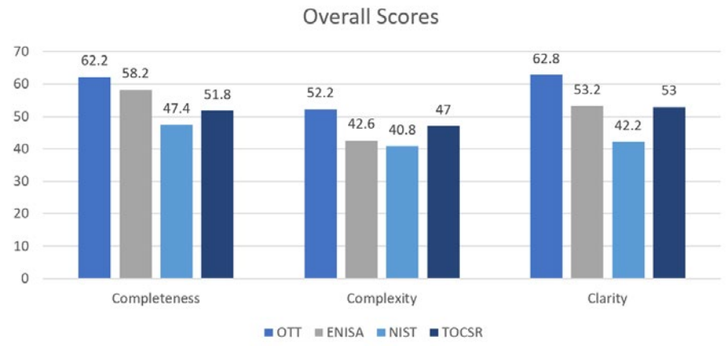

Figure 7: *Overall scores of each taxonomy by traits.*

evaluation because they are ubiquitous, descriptive words and encompass the individual characteristics that make a taxonomy a useful tool for communication. Therefore, respondents did not receive definitions for the traits. The rating score for each of these characteristics consists of a weighted scale from 1 to 5, from worst to best, with the following common descriptions: Not at all, Slightly, Moderately, Quite, Extreme. The weighted answers provide a quick method for scoring and comparing the taxonomies.

A consistent analysis method compares results for each of the traits without favoring one over another. However, organizations may choose to favor one trait over another because of its available resources. An organization may find the clarity of threat terms more advantageous than completeness, for example, if there is no intranet website for sharing a central glossary and training employees is unlikely. On the other hand, favoring clarity may also imply favoring the least complex taxonomy, and vice versa, given the relationship between these two traits.

Completeness

A complete threat taxonomy would be able to characterize all possible threat actions or events. The categories chosen by a taxonomy may preclude certain types of threats. For example, the NIST non-adversarial categories do not incorporate threats from legal action. For each taxonomy, respondents were asked to select one rating for the completeness of the taxonomy from these answers (with weight): Not at all complete (1), Slightly complete (2), Moderately complete (3), Quite complete (4), or Extremely complete (5). The score calculation is the average sum of weighted responses for each group. Therefore, groups with higher values in Figure 4 indicate more responses and the completeness scores in each group rank each taxonomy.

The overall completeness scores indicate OTT is the most complete. However, the Management group scores ENISA as the most complete. ENISA's taxonomy has the most threat actions present in the survey. Therefore, respondents may have given higher scores to ENISA based on this overall number. This is the most likely conclusion because respondents expect surveys to be brief. The low scores given to NIST further support this conclusion. NIST has the lowest number of threat actions in the survey because the length of adversarial threat actions in NIST SP 800-30 prevented listing them all in the survey application. The threat descriptions in NIST's adversarial tier create a cumbersome taxonomy table that is many pages long. The taxonomy review in previous sections shows that both NIST and OTT were unable to categorize legal threats. Additionally, NIST lacks more nuance for non-adversarial threats found in the other taxonomies. Even though scores for the TOCSR rank it third overall for completeness, the review in an earlier section did not find any events unfit for its threat categories. The business-risk perspective was likely a factor in lower completeness scores given its unique viewpoint from actions or failures of people, process, technology, or externalities.

Complexity

A complex threat taxonomy is one that is difficult to understand without additional context. Complexity could refer to either the structure or terms. Respondents may consider a taxonomy more complex if it has many high-level categories or more threat terms describing an event. For each taxonomy, respondents were asked to rate the overall complexity of each from these answers (with weight): Not at all complex (5), Slightly complex (4), Moderately complex (3), Quite complex (2), or Extremely complex (1). Score calculations follow the same process as in the completeness section. However, reversal of the weights is necessary to designate less complexity as the more desirable trait. Therefore, taxonomies with higher scores in Figure 5 are less complex. Respondents score the OTT and TOCSR as the least complex taxonomies. These taxonomies both have four top-tier categories with the most concise terminology to describe threat actions. The Financial Company and Management groups score TOCSR complexity just below the OTT. These groups are more likely to have a business-centric perspective that contributes to rating TOCSR higher than the other groups. However, these groups still rate the OTT as the least complex. Along with the Non-Management and Non-Financial groups both rating the OTT as the least complex, by larger margins, the overall score makes it the least complex taxonomy.

Clarity

A clear taxonomy would have simple threat terms and threat events that are logically relevant under the same category. While definitions are an essential element of a taxonomy for maintaining consistency, simple threat terms should plainly characterize a common set of threat events. For each taxonomy, respondents were asked to select a

rating for the clarity of terms from these answers (with weight): Not at all clear (1), Slightly clear (2), Moderately clear (3), Quite clear (4), or Extremely clear (5). The score calculation is the average sum of weighted responses for each group. Thus, the clearest taxonomies in Figure 6 have a higher score.

All the respondent groups rate the OTT as the clearest taxonomy. Only in the Management group did both the ENISA and TOCSR taxonomies have clarity scores close to the OTT. The respondent groups rate ENISA second, or a close third, in clarity. High clarity scores for ENISA's taxonomy were unexpected because of its alternative terms for several categories. Although, respondents may have seen the alternative terms as more descriptive characteristics for a category. Even though the TOCSR has the most concise terms for threat actions, its business-risk perspective appears to have detracted from the overall understanding of the terms by respondents.

Overall

The Open Threat Taxonomy overall scores are the highest for the completeness, complexity, and clarity traits. The combined group scores for each trait are viewable side-by-side in Figure 7. While the overall preference is for the OTT, both ENISA and TOCSR have strengths in different traits. The TOCSR has a high score for complexity, and the completeness score for ENISA is high. An organization favoring complexity or completeness may also consider either of these taxonomies. However, when it comes to clarity, the OTT outscores the other taxonomies by a large margin of at least ten points.

CONCLUSION

Survey respondents were asked to rate the clarity of terms to determine which threat taxonomy had the simplest terms and most logical grouping. Simple terms can help an organization's leadership understand threats to operations dependent on information technology. Many threat terms are available in CTI standards for intrusion analysis. However, there are too many terms for non-technical decision-makers to understand. Additionally, threat categories that logically group similar terms are clearer.

Review of the structure and terms of each threat taxonomy in the survey allowed respondents to judge which is the least complex. The exhaustive detail and multiple relationships within CTI standards that make them good for intrusion analysis also make them a poor choice for communicating with leadership. A smaller set of threat categories can reduce the complexity of cyberattacks for this audience. Grouping threat events with a hierarchical system can also reduce complexity when each category has similar events. The multiple levels within a hierarchical taxonomy provide several granularity options. This structure allows an organization to use the appropriate level for its risk assessment as it matures. Higher levels can help keep the risk assessment simple and small when it is immature. Lower levels can provide greater detail for complex threat scenarios when the organization is ready.

In order to assess the degree of inclusiveness for each threat taxonomy, the survey inquired about the completeness. Cyberattacks are not the only threats to an organization's information technology. Threats may arise from natural disasters, legal discussions or political interests, or employee accidents. The CTI standards concentrate on an adversary's malicious activity, so the lexicon in these standards is missing terms that characterize alternative threat sources. Risk frameworks help model all types of threats facing an organization. Comprehensive threat taxonomies fit into risk assessments, like NIST SP 800-30, to present decision-makers with a risk comparison across all of the threats.

This research found several methods for categorizing all of the possible threats to information technology. Only a handful of these threat taxonomies attempted to address all potential threats to IT within an organization. These nascent threat taxonomies may not be inclusive of all possible threats. The most mature taxonomy is about eight years old and updates have been infrequent. Since threat actions are one of the primary inputs for assessing IT risk, a public consensus of all the threats to information technology can improve communication within and between organizations.

The evaluation by both management and non-management personnel of these threat taxonomies strengthens the results of this research. The opinions of these two groups are vital for different reasons. Management needs to understand threats to improve communications with analysts and other business units in order to make quick decisions that influence the security resources of an organization. Non-management needs to present the threats to management, so they might obtain the necessary resources to address increasing threats. A familiar set of threat terms in meetings, reports, metrics, and risk assessments can help improve this communication. Based on the rating given for completeness, complexity, and clarity, this evaluation suggests each group prefers the Open Threat Taxonomy. This threat taxonomy can provide a complete picture of threat actions, with clear terms, in a manner that is simple for an organization's leadership to understand.

Future Research

This analysis resulted in the selection of a preferred threat taxonomy. However, this evaluation excludes an assessment of taxonomies to aid in decision-making by leadership. Evaluation of decision-making would require implementing a taxonomy into a risk assessment, mapping to security controls, and reviewing the issues which may arise from this implementation. Many of the risk frameworks present qualitative methods for assessments, but a quantitative assessment may favor one taxonomy over another. A comparative case study utilizing different threat taxonomies for threat scenarios with different risk frameworks, or the same risk framework with different assessment techniques are two possible evaluation ideas. Keys to success for this implementation would include mapping to security controls, like NIST SP 800-53, or security requirements, like NIST SP 800-171, and calculating probabilities of occurrence and impact based on changes to the threat landscape.

REFERENCES

[1] Definition of Virus. (2018, February 11). Retrieved from https://www.medicinenet.com/script/main/art.asp?articlekey=5997

[2] Glossary of Security Terms. (2018, February 11). Retrieved from https://www.sans.org/security-resources/glossary-of-terms/

[3] US Department of Defense, Joint Chiefs of Staff. (2013, October 22). Joint Intelligence (Joint Publication (JP) 2-0).

[4] Kime, B. P. (2016). Threat Intelligence: Planning and Direction. SANS Institute. Retrieved from https://www.sans.org/reading-room/whitepapers/threats/threat-intelligence-planning-direction-36857

[5] Lee, R. M., Cloppert, M. (2016). Forensics 578: Cyber Threat Intelligence. The SANS Institute. www.sans.org/course/cyber-threat-intelligence

[6] Farnham, G. (2013). Tools and Standards for Cyber Threat Intelligence Projects. SANS Institute. Retrieved from https://www.sans.org/reading-room/whitepapers/warfare/tools-standards-cyber-threat-intelligence-projects-34375

[7] Shackleford, D. (2017). Cyber Threat Intelligence Uses, Successes and Failures: The SANS 2017 CTI Survey. SANS Institute. Retrieved from www.sans.org/reading-room/whitepapers/analyst/cyber-threat-intelligence-uses-successes-failures-2017-cti-survey-37677

[8] MITRE Corporation. (2017, December 8). CAPEC - Common Attack Pattern Enumeration and Classification (CAPEC). Retrieved from https://capec.mitre.org/

[9] Mattern, T., Felker, J., Borum, R., Bamford, G. (2014). Operational Levels of Cyber Intelligence. International Journal of Intelligence and CounterIntelligence. Published Aug 6, 2014. Volume 27, Issue 4, Pages 702-719. DOI: 10.1080/08850607.2014.924811

[10] Shackleford, D. (2015). Who's Using Cyberthreat Intelligence and How? SANS Institute. Retrieved from https://www.sans.org/reading-room/whitepapers/analyst/cyberthreat-intelligence-how-35767

[11] Shackleford, D. (2016). The SANS State of Cyber Threat Intelligence Survey: CTI Important and Maturing. SANS Institute. Retrieved from https://www.sans.org/reading-room/whitepapers/bestprac/state-cyber-threat-intelligence-survey-cti-important-maturing-37177

[12] National Institute of Standards and Technology. (2010). Guide for Applying the Risk Management Framework to Federal Information Systems: a Security Life Cycle Approach. Special Publication (NIST SP) 800-37 Rev 1. http://dx.doi.org/10.6028/NIST.SP.800-37r1

[13] National Institute of Standards and Technology. (2012). Guide for Conducting Risk Assessments. Special Publication (NIST SP) 800-30 Rev 1. http://dx.doi.org/10.6028/NIST.SP.800-30r1

[14] Tarala, J., Tarala, K. (2015). Open Threat Taxonomy version 1.1. Enclave Security. Retrieved from http://www.auditscripts.com/resources/open_threat_taxonomy_v1.1a.pdf

[15] European Union Agency for Network and Information Security. (2016). ENISA Threat Landscape. ENISA. https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends/enisa-threat-landscape/enisa-threat-landscape

[16] Hutchins, E. M., Cloppert, M. J., Amin, R. M. (n.d.). Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. Lockheed Martin Corporation (Lockheed). Retrieved from https://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

[17] Cebula, James., Popeck, Mary., & Young, Lisa. (2014). A Taxonomy of Operational Cyber Security Risks Version 2 (CMU/SEI-2014-TN-006). Retrieved from the Software Engineering Institute, Carnegie Mellon University website: http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=91013

[18] Department of Homeland Security Risk Steering Committee. (2008, September) DHS Risk Lexicon. Department of Homeland Security (DHS). http://www.dhs.gov/xlibrary/assets/dhs_risk_lexicon.pdf

[19] Office of the Director of National Intelligence (ODNI). (n.d.). Cyber Threat Framework Version 4.0 Lexicon of concepts and definitions. Retrieved July 15, 2017 from https://www.dni.gov/files/ODNI/documents/features/Cyber_Threat_Framework_Lexicon.pdf

[20] Jakobsson, M. (2017). The Threat Taxonomy: A Working Framework to Describe Cyber Attacks. Agari. Retrieved on October 28, 2017 from https://www.agari.com/threat-taxonomy-framework-cyber-attacks/

[21] Burger, E.W., Goodman, M. D., Kampanakis, P., Zhu K. A. (2014). Taxonomy Model for Cyber Threat Intelligence Information Exchange Technologies. Proceedings of the 2014 ACM Workshop on Information Sharing & Collaborative Security. Pages 51-60. DOI: 10.1145/2663876.2663883

[22] Coburn, A.W.; Bowman, G.; Ruffle, S.J.; Foulser-Piggott, R.; Ralph, D.; Tuveson, M.; (2014). A Taxonomy of Threats for Complex Risk Management. Cambridge Risk Framework series. Centre for Risk Studies, University of Cambridge.

ABOUT THE AUTHOR

STEVEN LAUNIUS has over 10 years of professional information security experience in the financial services industry. His childhood passion for computers and puzzle games grew into a technology career to include systems support, audit, penetration testing, security consulting, and he currently manages a cyber-enabled fraud intelligence team at Discover Financial Services. Steven is a candidate in the Master of Science in Information Security Engineering program at SANS Technology Institute and holds the CISSP, CISA, and ten GIAC certifications.

APPENDIX A: THREAT TAXONOMY DETAILS

Note: The content as presented below was edited for presentation in the research survey; see references for complete taxonomies with definitions.

ENISA Threat Taxonomy

Physical attack (deliberate/ intentional)	
» Fraud » Sabotage » Vandalism » Theft (devices, storage media and documents) » Information leakage/sharing	» Unauthorized physical access / Unauthorized entry to premises » Coercion, extortion or corruption » Damage from the warfare » Terrorists attack
Unintentional damage / loss of information or IT assets	
» Information leakage/sharing due to human error » Erroneous use or administration of devices and systems » Using information from an unreliable source » Unintentional change of data in an information system » Inadequate design and planning or improperly adaptation	» Damage caused by a third party » Damages resulting from penetration testing » Loss of information in the cloud » Loss of (integrity of) sensitive information » Loss of devices, storage media and documents » Destruction of records
Disaster (natural, environmental)	
» Disaster (natural earthquakes, floods, landslides, tsunamis, heavy rains, heavy snowfalls, heavy winds) » Fire » Pollution, dust, corrosion » Thunder stroke » Water	» Explosion » Dangerous radiation leak » Unfavorable climatic conditions » Major events in the environment » Threats from space / Electromagnetic storm » Wildlife
Failures/ Malfunction	
» Failure of devices or systems » Failure or disruption of communication links (communication networks) » Failure or disruption of main supply	» Failure or disruption of service providers (supply chain) » Malfunction of equipment (devices or systems)
Outages	
» Loss of resources » Absence of personnel » Strike	» Loss of support services » Internet outage » Network outage
Eavesdropping/ Interception/ Hijacking	
» War driving » Intercepting compromising emissions » Interception of information » Interfering radiation » Replay of messages	» Network Reconnaissance, Network traffic manipulation and Information gathering » Man in the middle/ Session hijacking
Nefarious Activity/ Abuse	
» Identity theft (Identity Fraud/ Account) » Receive of unsolicited E-mail » Denial of service » Malicious code/ software/ activity » Social Engineering » Abuse of Information Leakage » Generation and use of rogue certificates » Manipulation of hardware and software » Manipulation of information » Misuse of audit tools	» Misuse of information/ information systems (including mobile apps) » Unauthorized activities » Unauthorized installation of software » Compromising confidential information (data breaches) » Hoax » Remote activity (execution) » Targeted attacks (APTs etc.) » Failed of bussines process » Brute force » Abuse of authorizations
Legal	
» Violation of laws or regulations / Breach of legislation » Failure to meet contractual requirements	» Unauthorized use of IPR protected resources » Abuse of personal data » Judiciary decisions/court orders

OTT Threat Actions & Ratings

Physical Threats	
» Loss of Property » Theft of Property » Accidental Destruction of Property » Natural Destruction of Property » Intentional Destruction of Property » Intentional Sabotage of Property » Intentional Vandalism of Property	» Electrical System Failure » HVAC Failure » Structural Facility Failure » Water Distribution System Failure » Sanitation System Failure » Natural Gas Distribution Failure » Electronic Media Failure
Resource Threats	
» Disruption of Water Resources » Disruption of Fuel Resources » Disruption of Materials Resources » Disruption of Electrical Resources » Disruption of Transportation Services » Disruption of Communications Services	» Disruption of Emergency Services Disruption of Governmental Services » Supplier Viability » Supplier Supply Chain Failure » Logistics Provider Failures » Logistics Route Disruptions » Technology Services Manipulation
Personnel Threats	
» Personnel Labor / Skills Shortage » Loss of Personnel Resources » Social Engineering of Personnel Resources	» Disruption of Personnel Resources » Negligent Personnel Resources » Personnel Mistakes / Errors » Personnel Inaction
Technical Threats	
» Organizational Fingerprinting via Open Sources » System Fingerprinting » Credential Discovery » Misuse of System Credentials » Escalation of Privilege » Abuse of System Privileges » Memory Manipulation » Cache Poisoning » Physical Manipulation of Technical Device » Manipulation of Trusted System	» Cryptanalysis » Data Leakage / Theft » Denial of Service » Maintaining System Persistence » Manipulation of Data in Transit / Use » Capture of Data in Transit / Use » Replay of Data in Transit / Use » Misdelivery of Data » Capture of Stored Data » Manipulation of Stored Data » Application Exploitation

NIST Risk Assessment Threat Event Taxonomy Exemplary

Adversarial	
Perform reconnaissance and gather information » 5 sub-elements **Craft or create attack tools** » 6 sub-elements **Deliver/insert/install malicious capabilities** » 14 sub-elements **Exploit and compromise** » 17 sub-elements	**Conduct an attack** » 21 sub-elements **Achieve results** » 13 sub-elements **Maintain a presence or set of capabilities** » 2 sub-elements **Coordinate a campaign** » 6 sub-elements
Non-Adversarial	
» Spill sensitive information » Mishandling of critical and/ or sensitive information by authorized users » Incorrect privilege settings » Communications contention » Unreadable display » Earthquake » Fire	» Flood » Hurricane » Resource depletion » Introduction of vulnerabilities into software products » Disk error » Pervasive disk error » Windstorm/tornado

Taxonomy of Operational Cyber Security Risks

Actions of People	
Inadvertent » Mistakes » Errors » Omissions **Deliberate** » Fraud » Sabotage » Theft » Vandalism	**Inaction** » Skills » Knowledge » Guidance » Availability
Systems and Technology Failures	
Hardware » Capacity » Performance » Maintenance » Obsolescence **Systems** » Design » Specifications » Integration » Complexity	**Software** » Compatibility » Configuration management » Change control » Security settings » Coding practices » Testing
Failed Internal Processes	
Process controls » Status monitoring » Metrics » Periodic review » Process ownership **Supporting Processes** » Staffing » Funding » Training and development » Procurement	**Process design or execution** » Process flow » Process documentation » Roles and responsibilities » Notifications and alerts » Information flow » Escalation of issues » Service level agreements » Task hand-off
External Events	
Disasters » Weather event » Fire » Flood » Earthquake » Unrest » Pandemic **Legal issues** » Regulatory compliance » Legislation » Litigation	**Business issues** » Supplier failure » Market conditions » Economic conditions **Service dependencies** » Utilities » Emergency services » Fuel » Transportation

TIMES CHANGE AND YOUR TRAINING DATA SHOULD TOO:

The Effect of Training Data Recency on Twitter Classifiers

By : **Ryan J O'Grady**, SANS Institute, Senior DevOps Engineer; Candidate, SANS Technology Institute, MS in Information Security Engineering

SOPHISTICATED ADVERSARIES ARE MOVING THEIR BOTNET COMMAND AND CONTROL INFRASTRUCTURE TO SOCIAL MEDIA MICROBLOGGING SITES SUCH AS TWITTER. AS SECURITY PRACTITIONERS WORK TO IDENTIFY NEW METHODS FOR DETECTING AND DISRUPTING SUCH BOTNETS, INCLUDING MACHINE-LEARNING APPROACHES, WE MUST BETTER UNDERSTAND WHAT EFFECT TRAINING DATA RECENCY HAS ON CLASSIFIER PERFORMANCE.

This research investigates the performance of several binary classifiers and their ability to distinguish between non-verified and verified tweets as the offset between the age of the training data and test data changed. Classifiers were trained on three feature sets: tweet-only features, user-only features, and all features. Key findings show that classifiers perform best at +0 offset, feature importance changes over time, and more features are not necessarily better. Classifiers using user-only features performed best, with a mean Matthews correlation coefficient of 0.95 ± 0.04 at +0 offset, 0.58 ± 0.43 at -8 offset, and 0.51 ± 0.21 at +8 offset. The R2 values are 0.90, 0.34, and 0.26, respectively. Thus, the classifiers tested with +0 offset accounted for 56% to 64% more variance than those tested with −8 and +8 offset. These results suggest that classifier performance is sensitive to the recency of the training data relative to the test data. Further research is needed to replicate this experiment with botnet vs. non-botnet tweets to determine if similar classifier performance is possible and the degree to which performance is sensitive to training data recency.

Photo Credit: Deposit Photos/everythingposs

INTRODUCTION

Botnets are using increasingly sophisticated methods not only to communicate but to conceal the presence of the botnet communication. Covert command and control channels make it more difficult to detect and disrupt botnet communication, one of the most common methods for disabling a botnet. Can modern machine learning techniques identify social media messages (tweets) associated with covert botnet command and control traffic? And if so, to what extent is the performance of such classifiers dependent on having recent training data? I selected tweet- and user-specific features of tweets and trained a variety of binary classifiers to distinguish between non-verified and verified tweets, then measured their performance when predicting the non-verified vs. verified status of tweets from a range of years. Developing a better understanding of how classifiers can distinguish between "trusted" and "untrusted" classes of tweets will lead to better techniques for detecting and disrupting covert botnet command and control channels. Understanding the impact of training data recency will lead to the creation of more effective classifiers.

"Botnets are massive, distributed networks of bots or zombies, typically seen in the form of malware-infected hosts – that is, unwitting participants."

LITERATURE REVIEW

Botnets are massive, distributed networks of bots or zombies, typically seen in the form of malware-infected hosts – that is, unwitting participants (Bailey, Cooke, Jahanian, Xu, & Karir, 2009; Vania, Meniya, & Jethva, 2013). Botnets rely on a command and control (C&C) channel to receive, execute, and respond to commands from the botmaster (Bailey et al., 2009; Vania et al., 2013). Over time, botnets have become more sophisticated, employing cutting-edge techniques to maintain availability and evade detection. In the early days of botnets, circa 2000, botnets relied on Internet Relay Chat (IRC) for C&C. IRC afforded a centralized command structure, anonymity, one-to-one (private) communication, and one-to-many communication (Vania et al., 2013). System administrators responded by restricting and monitoring access to IRC. Botnets, in turn, moved to a peer-to-peer (P2P) C&C structure, in which there is no central server; bots instead received commands from trusted locations or peers (other bots) (Bailey et al., 2009; Vania et al., 2013). Detecting such P2P botnets is difficult, and became more so with the advent of fast-flux botnets. Traditional approaches to detecting P2P botnets focus on network analysis – classifying traffic based on endpoints, latency, frequency, synchronicity, packet size, and similar (Bailey et al., 2009). The next evolution in botnet C&C was to move to social media and microblogging platforms, such as Facebook and Twitter (Kartaltepe, Morales, Xu, & Sandhu, 2010; Rodríguez-Gómez, Maciá-Fernández, & García-Teodoro, 2013; Stamp, Singh, H. Toderici, & Ross, 2013). The advantages of this move are that such platforms are resilient (their business depends on it), high-volume, and accessed over HTTP/HTTPS.

Moving to social media effectively hides C&C traffic amongst the noise of legitimate traffic, making it impossible to block outright (Kartaltepe et al., 2010; Stamp et al., 2013). Instead, defenders are forced to distinguish the C&C traffic from the legitimate traffic. Jose Nozario, a research scientist with Arbor Networks, documented a naïve approach to C&C over Twitter (2009) in which the botmaster issued commands via tweet, with each tweet containing a base64 encoded bit.ly link. The links, in turn, contained base64 encoded executables. Such a C&C approach is obvious to anyone that is watching, as base64 encoded messages stand out from typical Twitter traffic. Stegobot uses image steganography to hide the commands in images, which it then posts to social media – Facebook in this case – to conceal the presence of the command (Nagaraja et al., 2011). In *Covert Botnet Command and Control Using Twitter*, Pantic and Husain proposed an approach to concealing the presence of command and control traffic by using noiseless steganography and encoding the commands in the metadata – message length in this case (Pantic & Husain, 2015).

Disrupting a botnet requires either disabling the botmaster, disabling the zombies, cleaning the botnet malware from the zombies, or interfering with the ability of the botmaster to communicate with the zombies (Gu, Zhang, & Lee, 2008). Ten years ago, interfering with the C&C channel could be as simple as blocking all outbound IRC traffic from a network. Botnet detection focused on host-based or network-based analysis techniques (Cooke, Jahanian, & McPherson, 2005). The use of a legitimate social media platform as a covert C&C channel has complicated this. Research has been conducted on detecting Twitter spam accounts (Hua & Zhang, 2013) and classifying accounts as either human, bot, or cyborg (human-assisted bots or bot-assisted humans) (Chu, Gianvecchio, Wang, & Jajodia, 2010) using machine learning. In *Detection of Stegobot: a covert social network botnet*, Natarajan, Sheen, and Anitha (2012) presented a method for detecting Stegobot activity by analyzing the entropy of cover images, but there is still a research gap in the detection of botnet C&C behavior in Twitter. This gap is due at least in part to an acute lack of high-quality labeled training data, but this will hopefully change as researchers continue to identify large botnets in the wild, such as the 350,000 node botnet discovered by researchers at University College in London (Echeverría & Zhou, 2017). But the eventual availability of more and higher quality data leads to another research question: how important is the recency of the training data relative to the behaviors of interest? In other words, is there an expiration date on training data?

RESEARCH METHOD

The experiment required collecting and processing approximately 3.9 million tweets. From those, 1,500 tweets (1,000 non-verified and 500 verified) were randomly sampled from each year, 2010 – 2018, to generate nine datasets. For each year, several classifiers were trained with 1,000 tweets randomly sampled from that year's dataset. Based on the analysis conducted, there was no measurable increase in performance when training with more than 1,000 tweets per year. Classifiers were then evaluated against each year individually. Test sets were generated by randomly sampling 300 tweets (200 non-verified and 100 verified). The same training and test datasets were used for each classifier. To evaluate classifiers against their own year, tweets were drawn from the 500 tweets not used to train the classifier. To evaluate classifiers against other year data, the tweets were drawn from the entire dataset. This approach guaranteed that classifiers were not trained and evaluated using the same tweets.

Data Collection

Twitter introduced the Verified designation in June 2009. The initial analysis determined that the percentage of non-verified vs. verified tweets is highly imbalanced. As of March 2018, the data collected for this experiment shows that only 0.08% of Twitter users are verified, and 5% of tweets are from verified users. To address this, the data were oversampled without replacement, with a final distribution of 67% non-verified tweets and 33% verified tweets. However, the Twitter API limits access to only the 3,200 most recent tweets for each user. This makes it challenging to collect sufficient data for prior years. For example, of the 1,633 users in the final dataset, only 37 had one or more tweets in 2010. Of the 3.9 million tweets collected, 2,456 are from 2010. This is why it was necessary to collect such a large number of tweets to end up with a comparatively small dataset.

Table 1: *JSON and CSV Data Fields*

JSON Field Name	CSV Field Name	Data Type	Default
id	id	uint64[1]	*(none)*
created_at	created_at	datetime	*(none)*
text	text	string	*(none)*
truncated	is_truncated	uint81	0
source	source	string	*'' (empty string)*
lang	lang	string	'und'
is_quote_status	is_quote	uint8[1]	0
in_reply_to_status_id	is_reply	uint8[1,2]	0
retweeted_status	is_retweet	uint8[1,2]	0
quote_count	quote_count	uint64[1,3]	0
reply_count	reply_count	uint64[1,3]	0
retweet_count	retweet_count	uint64[1]	0
favorite_count	favorite_count	uint64[1]	0
place	has_place	uint8[1,2]	0
coordinates	has_coordinates	uint8[1,2]	0
user.id	user_id	uint64[1]	*(none)*
user.created_at	user_created_at	datetime	*(none)*
user.name	user_name	string	*'' (empty string)*
user.screen_name	user_screen_name	string	*'' (empty string)*
user.location	user_location	string	*'' (empty string)*
user.description	user_description	string	*'' (empty string)*
user.url	user_url	string	*'' (empty string)*
user.verified	user_verified	uint8[1]	0
user.followers_count	user_followers_count	uint64[1]	0
user.friends_count	user_friends_count	uint64[1]	0
user.listed_count	user_listed_count	uint64[1]	0
user.statuses_count	user_statuses_count	uint64[1]	0
user.favourites_count	user_favorites_count	uint64[1]	0

1: NumPy data type
2: Converted to boolean value based on presence/absence of JSON field
3: Field value is always 0 with free API

Data Processing

The Twitter API returns *statuses* as JSON objects. These JSON objects were then converted to comma-separated value (CSV) files. Some fields, such as *retweeted_status*, were converted to a boolean representation based on whether the field was present in the JSON object. Boolean values were converted to an integer value 0 (false) or 1 (true). Table 1 summarizes the JSON fields, corresponding CSV fields, Python data type, and default value used (if any). Table 2 summarizes the additional fields that were derived from the fields in Table 1.

Classifier Training

After data processing, the entire set of 3.9 million tweets was split into nine disjoint sets based on the *created_at* field year: 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, and 2018. For each year, a subset was created by randomly sampling 1,000 non-verified tweets and 500 verified tweets. This subset was then randomly split into a training set and a validation set, comprising 67% and 33% of the subset respectively. For each training set, the scikit-learn *Imputer* preprocessor was fitted to *NaN* values using the *mean* strategy and the *RobustScaler* preprocessor was fitted with default parameters. Finally, each classifier was fitted to the training data using default parameters.

Table 2: *Derived Fields*

CSV Field Name	Data Type	Default	Notes
created_ts	int64[1]	*(none)*	*created_at / 1,000,000,000*
user_created_ts	int64[1]	*(none)*	*user_created_at / 1,000,000,000*
retrieved_ts	int64[1]	*(none)*	derived from JSON timestamp
text_length	int64[1]	*(none)*	characters in *text*
text_length_pct	float64[1]	*(none)*	*text_length / 140 before 11/17* *text_length / 280 after 11/17*
age	int64[1]	*(none)*	*retrieved_ts – created_ts*
user_age	int64[1]	*(none)*	*retrieved_ts – user_created_ts*
quote_tsc *(time-scaled count)*	float64[1]	0	*quote_count / age*
reply_tsc	float64[1]	0	*reply_count / age*
retweet_tsc	float64[1]	0	*retweet_count / age*
favorite_tsc	float64[1]	0	*favorite_count / age*
user_followers_tsc	float64[1]	0	*user_followers_count / user_age*
user_friends_tsc	float64[1]	0	*user_friends_count / user_age*
user_listed_tsc	float64[1]	0	*user_listed_count / user_age*
user_statuses_tsc	float64[1]	0	*user_statuses_count / user_age*
user_favorites_tsc	float64[1]	0	*user_favorites_count / user_age*
time_delta	float64[1]	*user_avg_delta*	time since previous status
user_avg_delta	float64[1]	*(none)*	mean of non-null *time_delta* values for user
source_clean	string	*'' (empty string)*	stripped HTML

1: NumPy data type

Classifier Evaluation

Initial experimentation showed that classifiers performed extremely well by predicting that all tweets are non-verified. While accurate, this finding was of limited value. If 95% of tweets are from non-verified users and 5% are from verified users, a classifier would achieve 95% accuracy by predicting that all tweets are non-verified. This necessitated changing the performance metric from accuracy to Matthews correlation coefficient (MCC). MCC was chosen as it incorporates true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). The formula is as follows:

$$\mathrm{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$$

MCC is considered a balanced measure that performs well with imbalanced data classes (Boughorbel, Jarray, & El-Anbari, 2017). In the scenario where 95% of tweets are from non-verified users and 5% are from verified users, predicting that all tweets are non-verified would result in an MCC value of 0 (no correlation).

Each trained classifier was evaluated against each year from 2010 to 2018. For each year, a test set was created by randomly sampling 200 non-verified tweets and 100 verified tweets from the full dataset for the year. The exception to this was when the training year and test year were equal, in which case the validation set was used as the test set. Each test set was transformed using the previously fitted Imputer and RobustScaler preprocessors. Finally, the test class (non-verified vs. verified) was predicted using the previously fitted classifier and the MCC was calculated.

Potential Shortcomings

While care was taken to avoid common mistakes such as training and testing using the same data, the nature of the Twitter API does present some challenges. Chief among these is a potential selection bias. Twitter does not provide a method to choose a random user ID, so random user IDs were selected from a sample of the live stream, which means that all tweets used in this experiment are from users that were active as of March 2018.

Second, the Twitter API only provides the 3,200 most recent Tweets for each user. This means that the more active a user is, the more recent the cutoff. Consequently, data from more distant years are more likely to be from less active users.

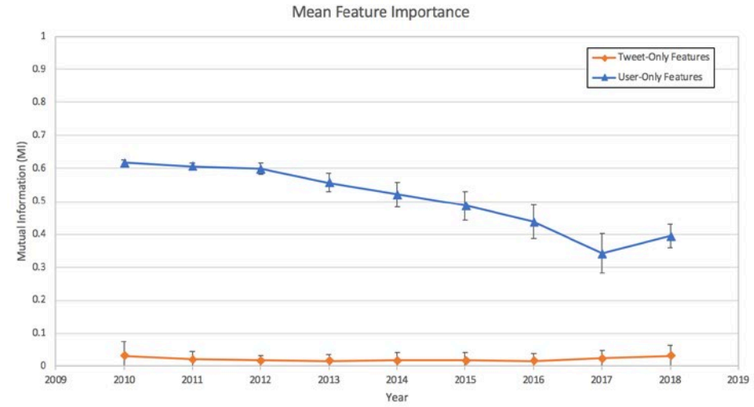

Figure 1: *Feature importance for tweet-only vs. user-only feature sets*

FINDINGS AND DISCUSSION

Feature Importance

Each classifier was trained and evaluated with three features sets: *tweet-only features*, *user-only features,* and *all features*.

Tweet-only features are specific to an individual tweet and do not include user information. The following features were included: created_ts, age, time_delta, is_truncated, is_quote, is_reply, is_retweet, has_place, has_coordinates, text_length, text_length_pct, quote_tsc, reply_tsc, retweet_tsc, and favorite_tsc.

User-only features are independent of individual tweets and represent a user at the time the tweet is retrieved, not at the time the tweet is originally posted. The following features were included: user_created_ts, user_age, user_followers_tsc, user_friends_tsc, user_listed_tsc, user_statuses_tsc, user_favorites_tsc, user_avg_delta.

All features combines the features from the tweet-only and user-only feature sets.

The mutual information (MI) score for each feature was calculated to determine how strongly the feature contributed to the correct classification. This information was not used for feature selection, but rather to measure the relative importance of features. Figure 1 is a plot of the mean and standard deviation per year of the tweet-only vs. user-only feature sets.

The mean mutual information scores for the user-only features are 14 to 20 times higher than those for tweet-only features. This means that user-only features are stronger indicators of whether a tweet is non-verified vs. verified, a conclusion that will be confirmed when classifier performances are examined.

The mean mutual information scores for user-only features drop steadily over time. One possible explanation for this observation is that user characteristics and behaviors have become less able to predict if a tweet is from a verified user as the Twitter user base has grown in size and diversity.

Year	2010	2011	2012	2013	2014	2015	2016	2017	2018
2010	1	0.97	0.746	0.511	0.379	0.222	0.149	0.071	0.071
2011	1	0.985	0.748	0.466	0.336	0.218	0.155	0.051	0.058
2012	0.993	0.97	0.978	0.748	0.603	0.423	0.267	0.102	0.053
2013	0.773	0.932	0.977	0.97	0.71	0.607	0.51	0.309	0.136
2014	0.762	0.933	0.911	0.887	0.895	0.711	0.679	0.473	0.2
2015	0.802	0.94	0.872	0.818	0.856	0.834	0.779	0.572	0.302
2016	0.41	0.813	0.86	0.772	0.858	0.742	0.864	0.682	0.312
2017	0.378	0.766	0.75	0.702	0.692	0.775	0.78	0.888	0.749
2018	-0.097	0.018	0.126	-0.071	0.237	0.372	0.541	0.788	0.857

Year	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8
2010									1	0.97	0.746	0.511	0.379	0.222	0.149	0.071	0.071
2011								1	0.985	0.748	0.466	0.336	0.218	0.155	0.051	0.058	
2012							0.993	0.97	0.978	0.748	0.603	0.423	0.267	0.102	0.053		
2013						0.773	0.932	0.977	0.97	0.71	0.607	0.51	0.309	0.136			
2014					0.762	0.933	0.911	0.887	0.895	0.711	0.679	0.473	0.2				
2015				0.802	0.94	0.872	0.818	0.856	0.834	0.779	0.572	0.302					
2016			0.41	0.813	0.86	0.772	0.858	0.742	0.864	0.682	0.312						
2017		0.378	0.766	0.75	0.702	0.692	0.775	0.78	0.888	0.749							
2018	-0.097	0.018	0.126	-0.071	0.237	0.372	0.541	0.788	0.857								

Figure 2: *Example of performance data before and after being offset*

Classifier Performance

The experiment compares the performance of seven classifiers: k-nearest neighbors, logistic regression, SVC, multi-layer perceptron, decision trees, random forest, and gradient boosting. In each instance, the scikit-learn implementation with default parameters was used. No optimization was performed.

For the purposes of evaluating classifier performance, the data has been offset such that the difference between training year and test year are the same at a given point on the x-axis. For example, Figure 2 shows MCC values for the SVC classifier with user-only features. In the top table, the y-axis represents the training data year and the x-axis represents the test data year. However, this research is primarily focused on determining how well classifiers can predict the class of data that is older or newer than the training data. In the bottom table, the data has been shifted to reflect the offset between the test year and the training year. In the first row, which corresponds to training data from 2010, the test data from 2010 has an offset of +0 (2010 – 2010), the test data from 2011 has an offset of +1 (2011 – 2010), and so on. As a result, we can quickly see that for an offset of +0 (that is, for all instances in which the training data year and test data year were the same), the performance was quite good, but declines as the offset grows. This approach to offsetting the data is used throughout this paper.

Tweet-Only Features

Figure 3 compares the individual mean Matthews correlation coefficient of seven different classifiers.

Figure 4 shows the mean Matthews correlation coefficient and standard deviation of all seven classifiers. Performance is optimal at +0 offset, with a mean MCC of 0.31 ± 0.04 (weak positive relationship), and quickly drops as the offset changes. After reaching a low of 0.09 ± 0.06 at +5 offset, the MCC begins to rise, but so does the standard deviation, until reaching an MCC of 0.2 ± 0.19.

User-Only Features

Figure 5 compares the individual mean Matthews correlation coefficient of seven different classifiers.

Figure 6 shows the mean Matthews correlation coefficient and standard deviation of all seven classifiers. Again, performance is optimal at +0 offset, with a mean MCC of 0.95 ± 0.04 (very strong positive relationship), and drops as the offset changes. Unlike the tweet-only features, however, performance at the extreme offsets still shows a moderate positive relationship. At –8 offset the mean MCC is 0.58 ± 0.43, while at +8 offset the mean MCC is 0.51 ± 0.21. In both instances, the standard deviation grows as we move from an offset of +0.

All Features

Figure 7 compares the individual mean Matthews correlation coefficient of seven different classifiers.

Figure 8 shows the mean Matthews correlation coefficient and standard deviation of all seven classifiers. Again, performance is optimal at +0 offset, with a mean MCC of 0.89 ± 0.08 (very strong positive relationship) and drops as the offset changes. Performance at the extreme offsets still shows a moderate positive relationship, though not to the same extent as with the user-only features. At −8 offset the mean MCC is 0.4 ± 0.41, while at +8 offset the mean MCC is 0.45 ± 0.3. In both instances, the standard deviation grows as we move from an offset of +0.

Summary of Findings

The experiment resulted in three key findings:

1. **Classifiers perform best at +0 year offset and generally perform worse as the offset increases or decreases.** Ensemble classifiers (random forest, gradient boosting), decision trees, and multi-layer perceptron were most resilient to this performance loss. For these four classifiers, the mean MCC at +0, –8, and +8 offset was 0.98 ± 0.01, 0.91 ± 0.08, and 0.51 ± 0.06 respectively. SVC was the most susceptible to performance loss, with mean MCC at +0, –8, and +8 offset of 0.92, –0.097, and 0.07 respectively. This finding suggests that it is important to use recent training data and that different classifiers are more or less sensitive to this recency requirement.
2. **Feature importance changes over time.** The tweet-only features were uniformly uninformative, with a mean mutual information (MI) ranging from 0.03 ± 0.04 in 2010 to 0.03 ± 0.03 in 2018. But the user-only features proved more informative, with a mean MI ranging from 0.62 ± 0.01 in 2010 to 0.39 ± 0.03. The mean MI change from 2010 to 2018 was –0.22, while the features with the greatest change were *user_friends_tsc* and *user_statuses_tsc* with MI changes of –0.32 and –0.24 respectively. The feature with the smallest change was *user_created_ts* with an MI change of –0.18. While this research does not attempt to identify what underlying user trends cause these changes in feature importance over time, it is important to know that they do change. This would explain, at least in part, the recency requirement from the first finding.
3. **Adding more features does not necessarily improve performance.** In fact, certain classifiers performed markedly worse with *all features* compared to *user-only features*. The mean MCC change was –0.13, while SVC and multi-layer perceptron both showed a mean MCC change of –0.26. The classifier with the smallest change (though still negative) was the random forest with a mean MCC change of –0.01. Gradient boosting and decision trees had a mean MCC change of –0.04 and –0.05 respectively. This finding throws into question a commonly-accepted machine learning maxim and merits further investigation.

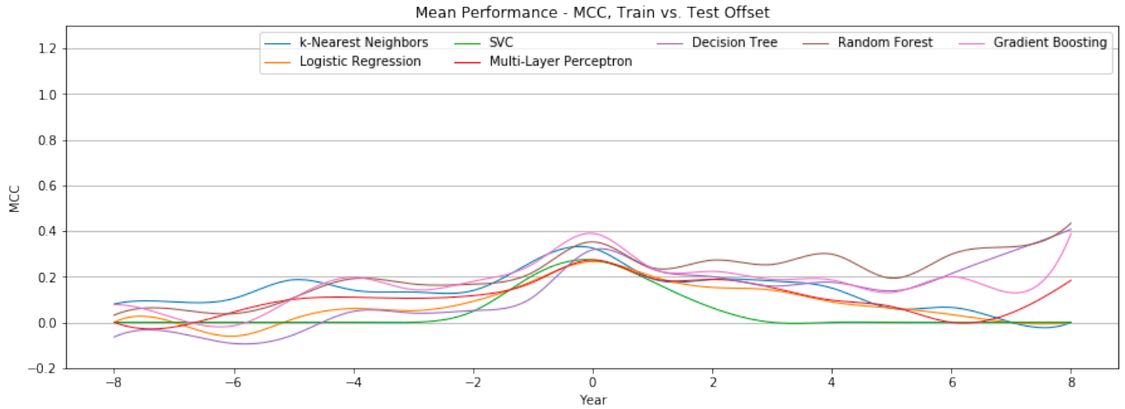

Figure 3: *Mean performance of each classifier with tweet-only features*

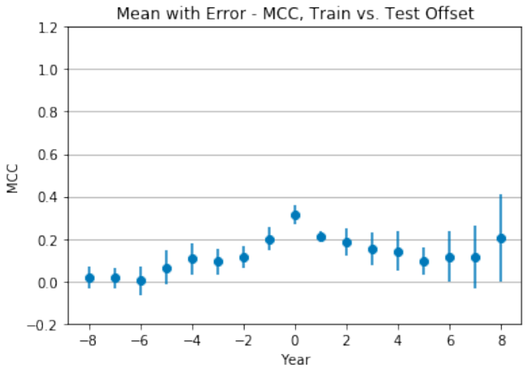

Figure 4: *Mean classification performance across all classifiers with tweet-only features*

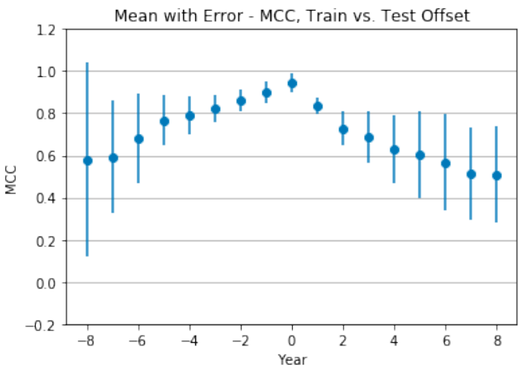

Figure 6: *Mean classification performance across all classifiers with user-only features*

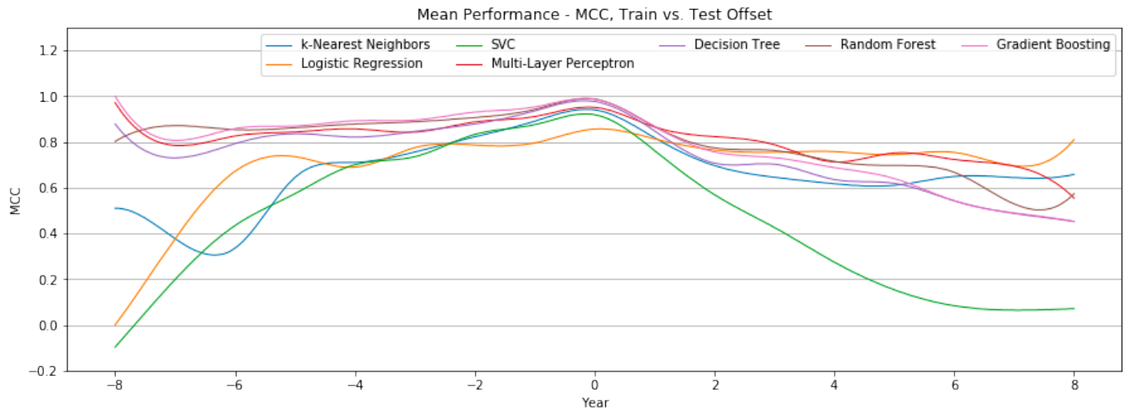

Figure 5: *Mean performance of each classifier with user-only features*

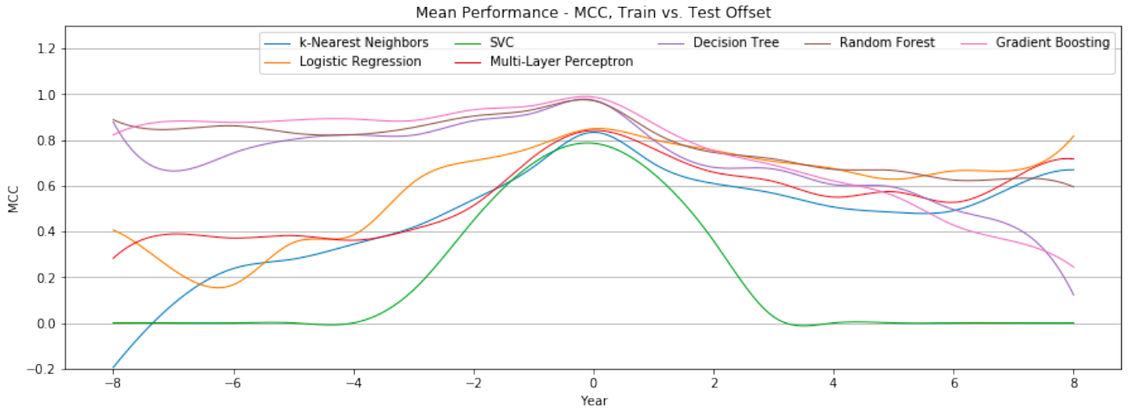

Figure 7: *Mean performance of each classifier with user-only features*

RECOMMENDATIONS AND IMPLICATIONS

Recommendations for Practice

This research has shown that ML classifiers can distinguish between non-verified and verified tweets and that such classification is highly sensitive to the recency of the training data. This work has significant implications for future research in detecting botnet traffic over Twitter. In particular, researchers must be sensitive to the timeliness of the training data they are using. Future research should test whether a similar effect is found when distinguishing between genuine traffic and botnet C&C traffic.

If the approach used in this research proves viable for botnet traffic as well, social media platforms (such as Twitter) could employ this approach to identify messages that are part of a botnet C&C channel and disrupt it by modifying/removing the messages. In addition, network owners could employ this approach at their network boundary, in conjunction with an SSL forward proxy, to inspect outbound messages to social media platforms and identify C&C messages. Such an approach would supplement host-based approaches to detection.

The drawbacks to such an approach are largely performance-cost related. Network owners would incur a relatively small processing cost by implementing such an approach on their network, given the relative volume of social media traffic compared to other traffic. However, getting buy-in from a significant number of network owners would be challenging. Conversely, social media platforms could implement such an approach to identifying suspicious traffic, but the processing cost incurred would be higher. The impact on processing speed for such an approach is outside the scope of this research. Lastly, as with any automated system, false positives have the potential to negatively affect users' experiences.

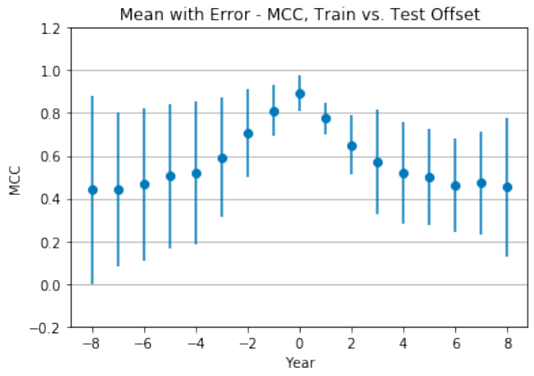

Figure 8: *Mean classification performance across all classifiers with all features*

Implications for Future Research

This is a field ripe for further exploration, and several significant questions directly follow from this research. First, and perhaps most obvious, does the methodology used in this research hold true for distinguishing between botnet C&C traffic and non-botnet C&C traffic? Related, is the assumption that there is a correlation between non-verified vs. verified and botnet vs. non-botnet tweets true? And if so, how strongly are they correlated?

Second, what effects do the limitations on the Twitter API have on the performance observed in this research? Would the ability to truly randomly select users and retrieve users' entire tweet history result in better or worse performance? Or to frame it differently, is the performance observed in this research a result of these limitations, or in spite of them?

Third, several features were ignored from the training data, including the actual tweet content. What effect would the inclusion of these features have on performance? Are there other features that could be derived from the available data to improve performance?

The ability to predict whether a tweet resembles tweets from verified or non-verified users does have some direct applicability. Twitter could use such an approach to streamline the Verified process and suggest that certain individuls apply for Verified status who strongly resemble other verified individuals. However, the ultimate goal of this research is to detect botnet traffic. As such, the obvious next step is to identify or create a labeled training set and repeat this research to answer the first question above.

CONCLUSION

This research addresses the question of whether ML classifiers can distinguish between non-verified and verified tweets, with those classes being stand-ins for botnet vs. non-botnet tweets. Further, the research investigates the degree to which the offset between training and test data affects the performance of the classifiers. In this research, I collected a corpus of tweets from 2010 ¬– 2018, trained a variety of classifiers on data from each year, and tested the performance of the classifiers against data from every year. This approach was repeated for three feature sets: tweet-only features, user-only features, and all features.

Findings showed that the classifiers could distinguish between non-verified and verified tweets, with the user-only feature set performing best. Further, the findings showed that there is a strong performance loss when testing a classifier against data from years other than the year used to train the classifier. In most cases, the performance loss becomes more severe as the gap between training and test year increases.

Future research is needed to identify what correlation (if any) there is between non-verified vs. verified and botnet vs. non-botnet tweets. In addition, future research is also needed to repeat this experiment with labeled botnet vs. non-botnet tweets.

ACKNOWLEDGEMENTS

I would like to acknowledge and thank the following people for providing guidance, feedback, and technical editing on this research effort: David Hoelzer, Chris MacLellan, Fernando Maymi, Caitlin Tenison, and Jack Zaientz.

REFERENCES

[1] Bailey, M., Cooke, E., Jahanian, F., Xu, Y., & Karir, M. (2009). A Survey of Botnet Technology and Defenses. In *2009 Cybersecurity Applications Technology Conference for Homeland Security* (pp. 299–304). https://doi.org/10.1109/CATCH.2009.40

[2] Boughorbel, S., Jarray, F., & El-Anbari, M. (2017). Optimal classifier for imbalanced data using Matthews Correlation Coefficient metric. *PLOS ONE*, *12*(6), e0177678. https://doi.org/10.1371/journal.pone.0177678

[3] Chu, Z., Gianvecchio, S., Wang, H., & Jajodia, S. (2010). Who is Tweeting on Twitter: Human, Bot, or Cyborg? In *Proceedings of the 26th Annual Computer Security Applications Conference* (pp. 21–30). New York, NY, USA: ACM. https://doi.org/10.1145/1920261.1920265

[4] Cooke, E., Jahanian, F., & McPherson, D. (2005). The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets. In *Proceedings of the Steps to Reducing Unwanted Traffic on the Internet on Steps to Reducing Unwanted Traffic on the Internet Workshop* (pp. 6–6). Berkeley, CA, USA: USENIX Association. Retrieved from http://dl.acm.org/citation.cfm?id=1251282.1251288

[5] Echeverría, J., & Zhou, S. (2017). Discovery, Retrieval, and Analysis of "Star Wars" botnet in Twitter. *ArXiv:1701.02405 [Cs]*. Retrieved from http://arxiv.org/abs/1701.02405

[6] Gu, G., Zhang, J., & Lee, W. (2008). *BotSniffer: Detecting Botnet Command and Control Channels in Network Traffic.*

[7] Hua, W., & Zhang, Y. (2013). Threshold and Associative Based Classification for Social Spam Profile Detection on Twitter. In *2013 Ninth International Conference on Semantics, Knowledge and Grids* (pp. 113–120). https://doi.org/10.1109/SKG.2013.15

[8] Jose Nazario. (2009, August 13). Twitter-based Botnet Command Channel. Retrieved May 27, 2018, from https://asert.arbornetworks.com/twitter-based-botnet-command-channel/

[9] Kartaltepe, E. J., Morales, J. A., Xu, S., & Sandhu, R. (2010). Social Network-Based Botnet Command-and-Control: Emerging Threats and Countermeasures. In J. Zhou & M. Yung (Eds.), *Applied Cryptography and Network Security* (pp. 511–528). Berlin, Heidelberg: Springer Berlin Heidelberg.

[10] Nagaraja, S., Houmansadr, A., Piyawongwisal, P., Singh, V., Agarwal, P., & Borisov, N. (2011). Stegobot: A Covert Social Network Botnet. In T. Filler, T. Pevný, S. Craver, & A. Ker (Eds.), *Information Hiding* (pp. 299–313). Berlin, Heidelberg: Springer Berlin Heidelberg.

[11] Natarajan, V., Sheen, S., & Anitha, R. (2012). Detection of StegoBot: A Covert Social Network Botnet. In *Proceedings of the First International Conference on Security of Internet of Things* (pp. 36–41). New York, NY, USA: ACM. https://doi.org/10.1145/2490428.2490433

[12] Pantic, N., & Husain, M. I. (2015). Covert Botnet Command and Control Using Twitter. In *Proceedings of the 31st Annual Computer Security Applications Conference* (pp. 171–180). New York, NY, USA: ACM. https://doi.org/10.1145/2818000.2818047

[13] Rodríguez-Gómez, R. A., Maciá-Fernández, G., & García-Teodoro, P. (2013). Survey and Taxonomy of Botnet Research Through Life-cycle. *ACM Comput. Surv.*, *45*(4), 45:1–45:33. https://doi.org/10.1145/2501654.2501659

[14] Stamp, M., Singh, A., H. Toderici, A., & Ross, K. (2013). Social Networking for Botnet Command and Control. *International Journal of Computer Network and Information Security*, *5*, 11–17.

[15] Vania, J., Meniya, A., & Jethva, H. (2013). A Review on Botnet and Detection Technique. *International Journal of Computer Trends and Technology*, *4*(1), 23–29.

ABOUT THE AUTHOR

RYAN J. O'GRADY has worked in the software and information security field for over 14 years. In 2012, he joined Soar Technology as the lead Research Scientist for Cyberspace Workforce Development. He was the Principal Investigator for an Air Force Research Labs (AFRL) project to develop an intelligent training system for cyberspace operators that enables individualized, personalized training in realistic environments. Mr. O'Grady was also the Technical Lead on a related project to create cognitive models of cyber attackers for training, testing, and evaluation purposes. Prior to that, he worked as a Software Architect for Army TACOM, where he oversaw the migration and security of a production enclave, and as a Research Engineer for Cybernet Systems, where he performed R&D for a variety of DoD customers. Mr. O'Grady earned his B.S.E. in Computer Science from the University of Michigan in 2004 and is pursuing a M.S. in Information Security Engineering from the SANS Technology Institute. Certifications: GCPM, GSEC, GCIH, GCIA, CPTE, Security+

APPENDIX A: CLASSIFIERS

This appendix lists the scikit-learn classifiers and parameters used in this research, as shown by a call to each classifiers' *str()* method: `NeighborsClassifier(algorithm='auto', leaf_size=30, metric='minkowski', metric_params=None, n_jobs=1, n_neighbors=5, p=2, weights='uniform')`

```
LogisticRegression(C=1.0, class_weight=None, dual=False, fit_intercept=True, intercept_
      scaling=1, max_iter=100, multi_class='ovr', n_jobs=1, penalty='l2', random_
      state=42, solver='liblinear', tol=0.0001, verbose=0, warm_start=False)

SVC(C=1.0, cache_size=200, class_weight=None, coef0=0.0, decision_function_shape='ovr',
      degree=3, gamma='auto', kernel='rbf', max_iter=-1, probability=False,
      random_state=42, shrinking=True, tol=0.001, verbose=False)

MLPClassifier(activation='relu', alpha=0.0001, batch_size='auto', beta_1=0.9,
      beta_2=0.999, early_stopping=False, epsilon=1e-08, hidden_layer_sizes=(100,),
      learning_rate='constant', learning_rate_init=0.001, max_iter=200, momentum=0.9,
      nesterovs_momentum=True, power_t=0.5, random_state=42, shuffle=True, solver='adam',
      tol=0.0001, validation_fraction=0.1, verbose=False, warm_start=False)

DecisionTreeClassifier(class_weight=None, criterion='gini', max_depth=None, max_
      features=None, max_leaf_nodes=None, min_impurity_decrease=0.0, min_
      impurity_split=None, min_samples_leaf=1, min_samples_split=2, min_weight_
      fraction_leaf=0.0, presort=False, random_state=42, splitter='best')

RandomForestClassifier(bootstrap=True, class_weight=None, criterion='gini', max_depth=None,
      max_features='auto', max_leaf_nodes=None, min_impurity_decrease=0.0, min_impurity_
      split=None, min_samples_leaf=1, min_samples_split=2, min_weight_fraction_leaf=0.0, n_
      estimators=10, n_jobs=1, oob_score=False, random_state=42, verbose=0, warm_start=False)

GradientBoostingClassifier(criterion='friedman_mse', init=None, learning_rate=0.1,
      loss='deviance', max_depth=3, max_features=None, max_leaf_nodes=None, min_
      impurity_decrease=0.0, min_impurity_split=None, min_samples_leaf=1, min_samples_
      split=2, min_weight_fraction_leaf=0.0, n_estimators=100, presort='auto',
      random_state=42, subsample=1.0, verbose=0, warm_start=False)
```


APPENDIX B: INDIVIDUAL CLASSIFIER PERFORMANCE

K-Nearest Neighbors

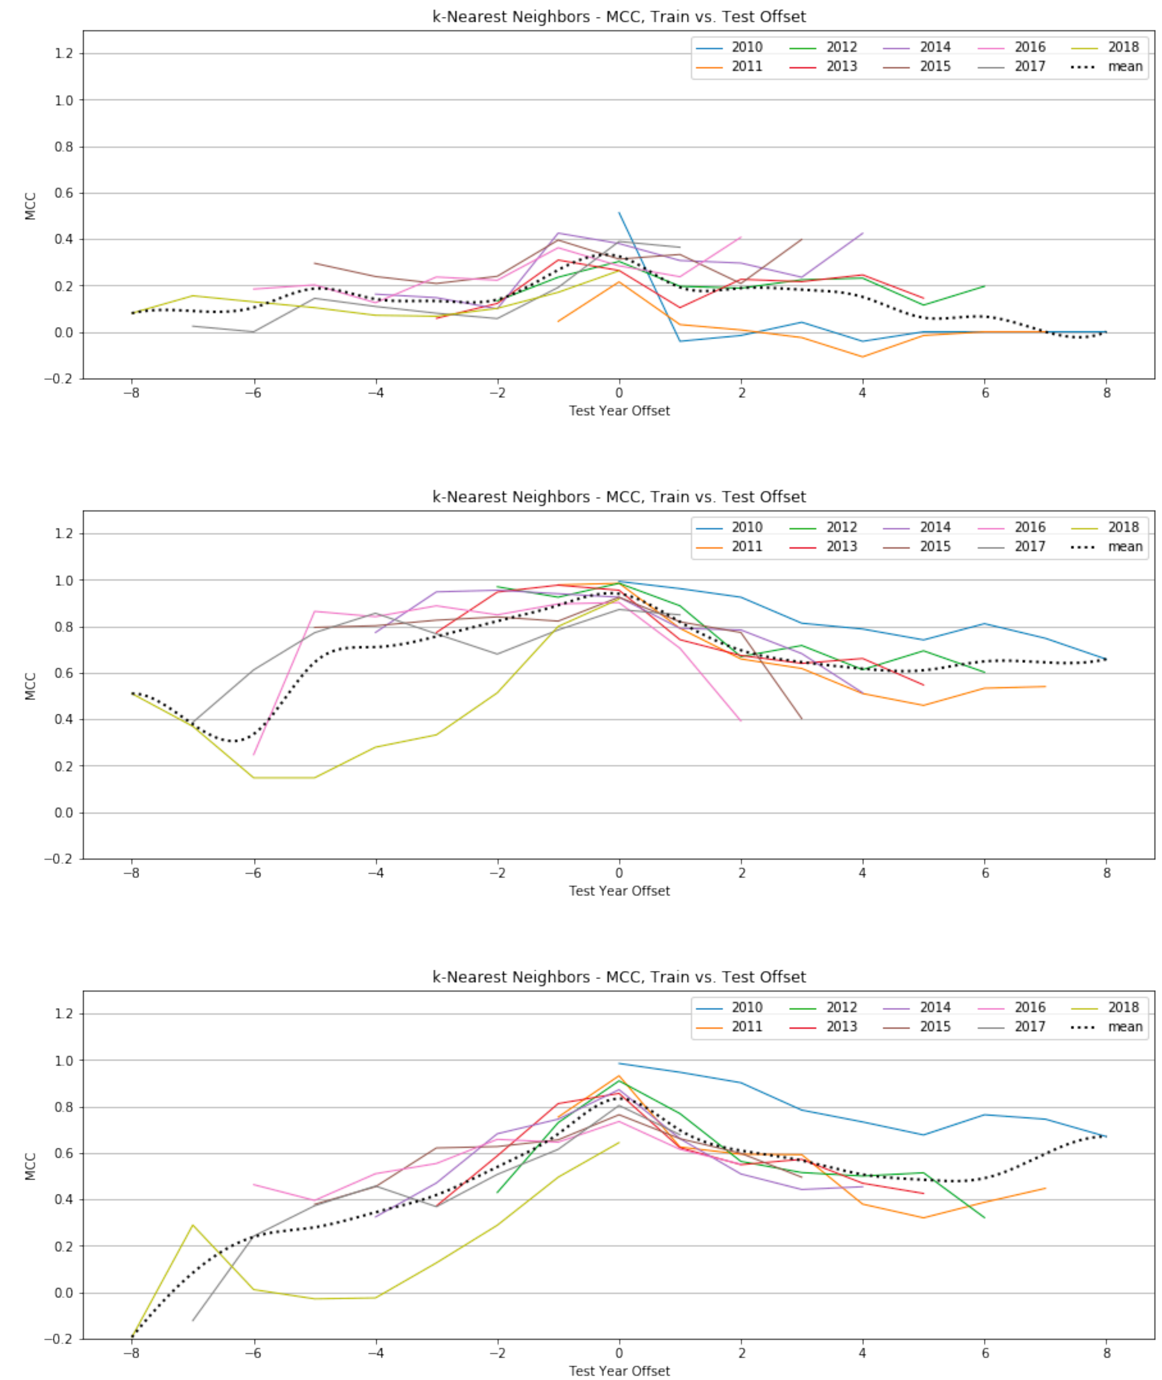

Figure 9: k-Nearest-Neighbors, MCC by year, Tweet-Only vs. User-Only vs. All Features

Logistic Regression

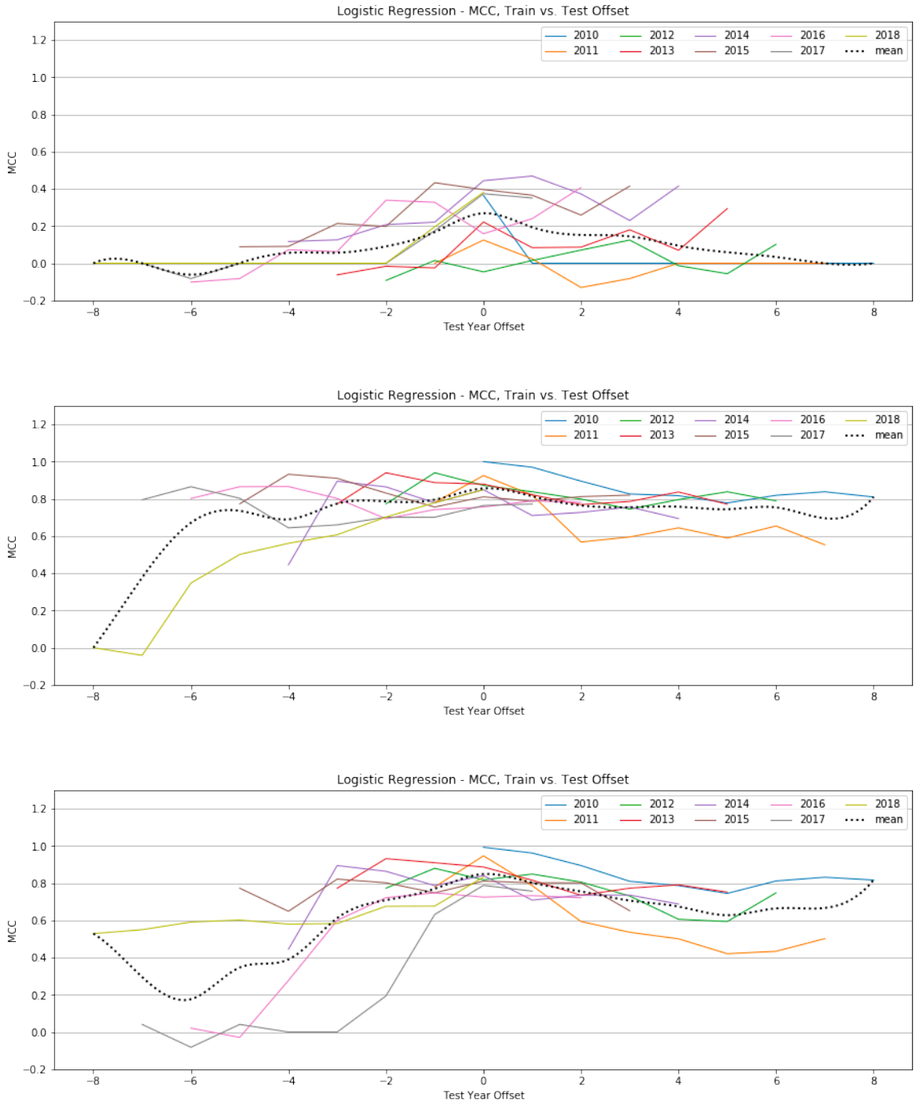

Figure 10: Logistic Regression, MCC by year, Tweet-Only vs. User-Only vs. All Features

SVC

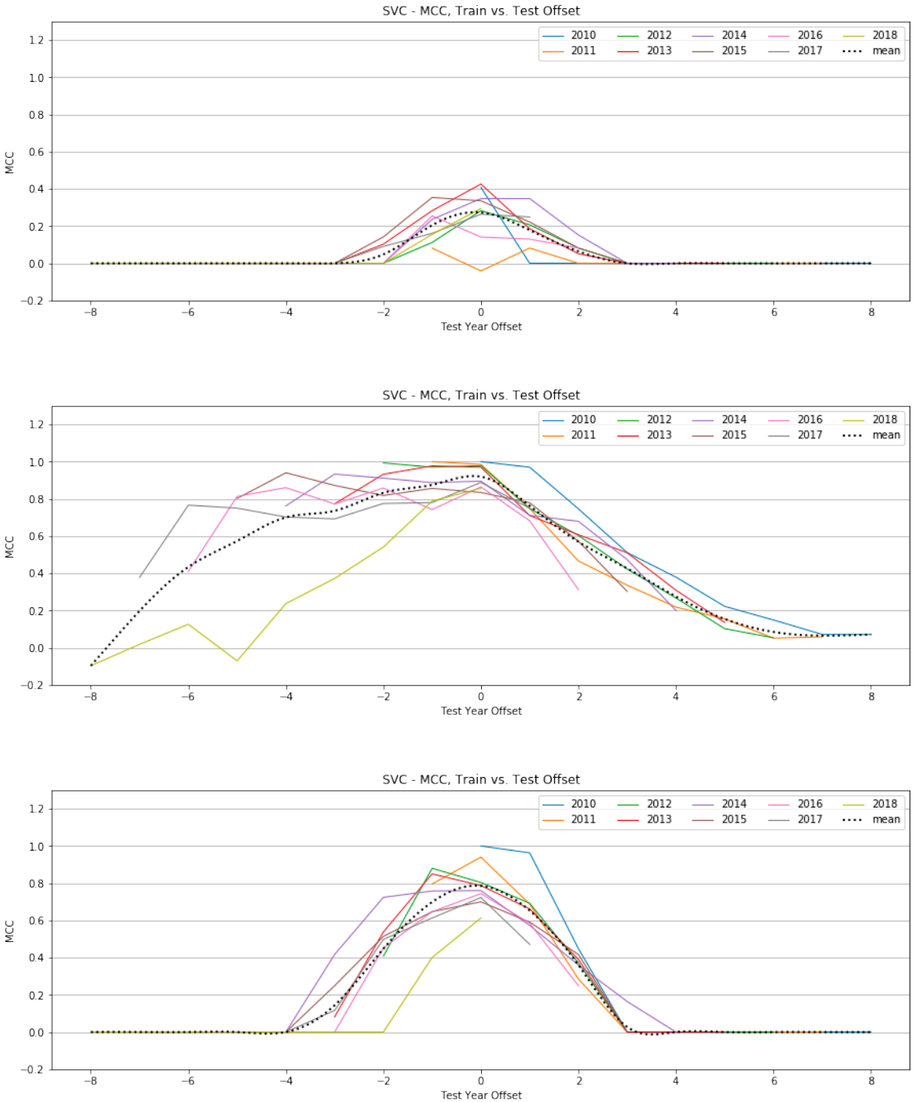

Figure 11: SVC, MCC by year, Tweet-Only vs. User-Only vs. All Features

MLP

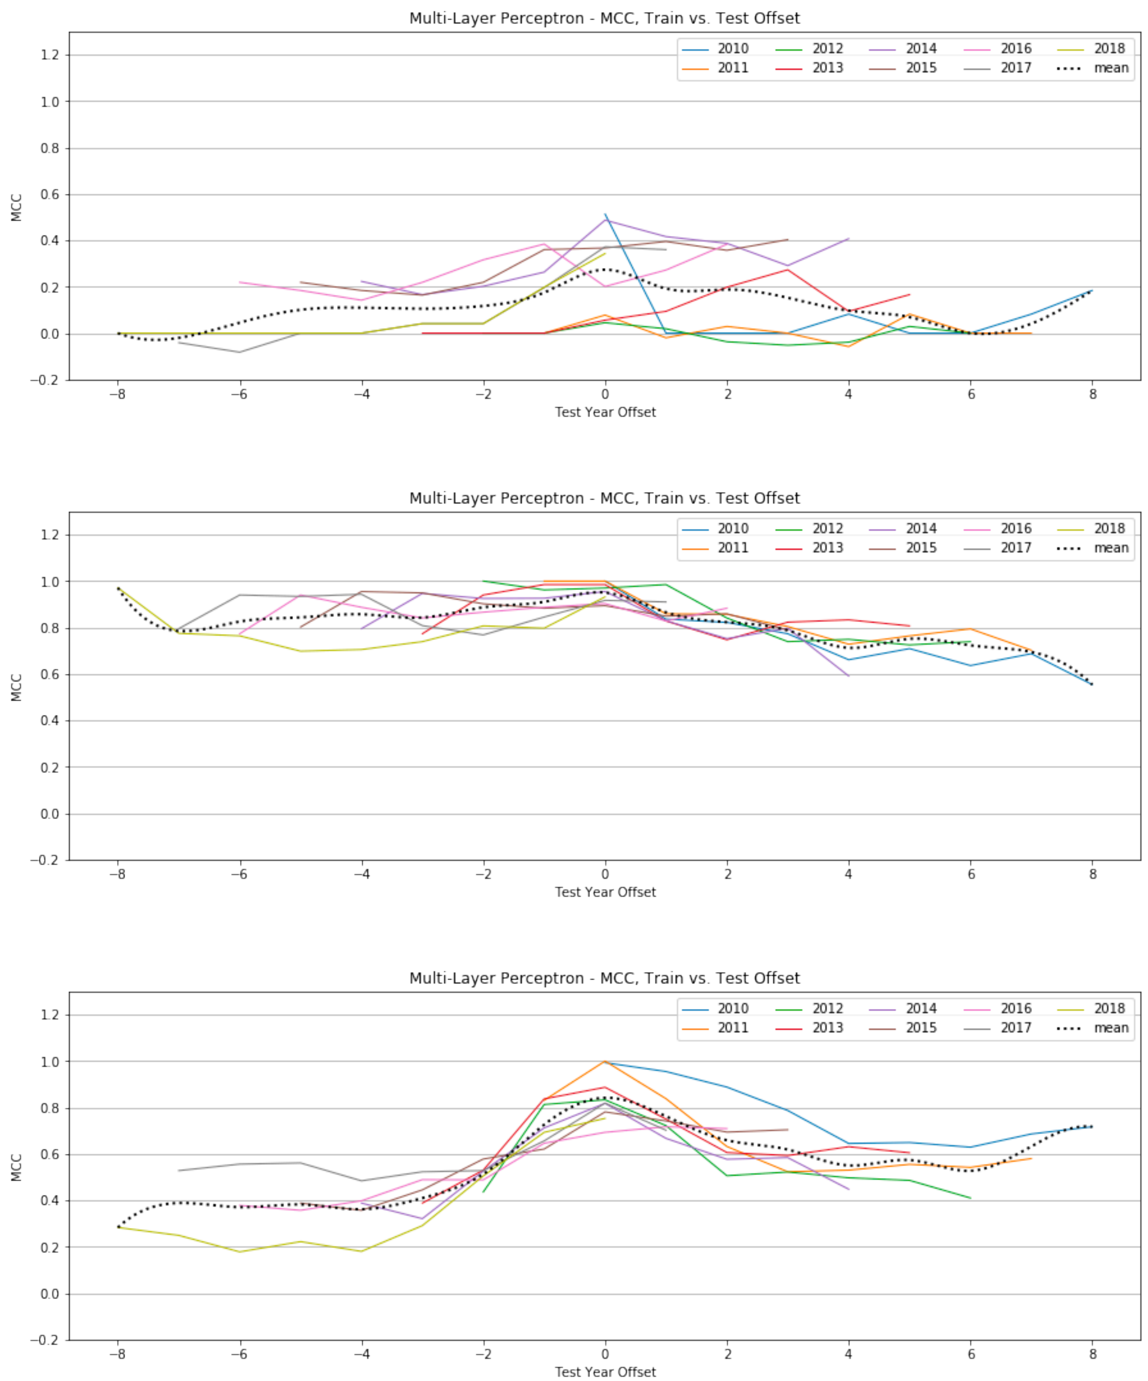

Figure 12: Multi-Layer Perceptron, MCC by year, Tweet-Only vs. User-Only vs. All Features

Decision Trees

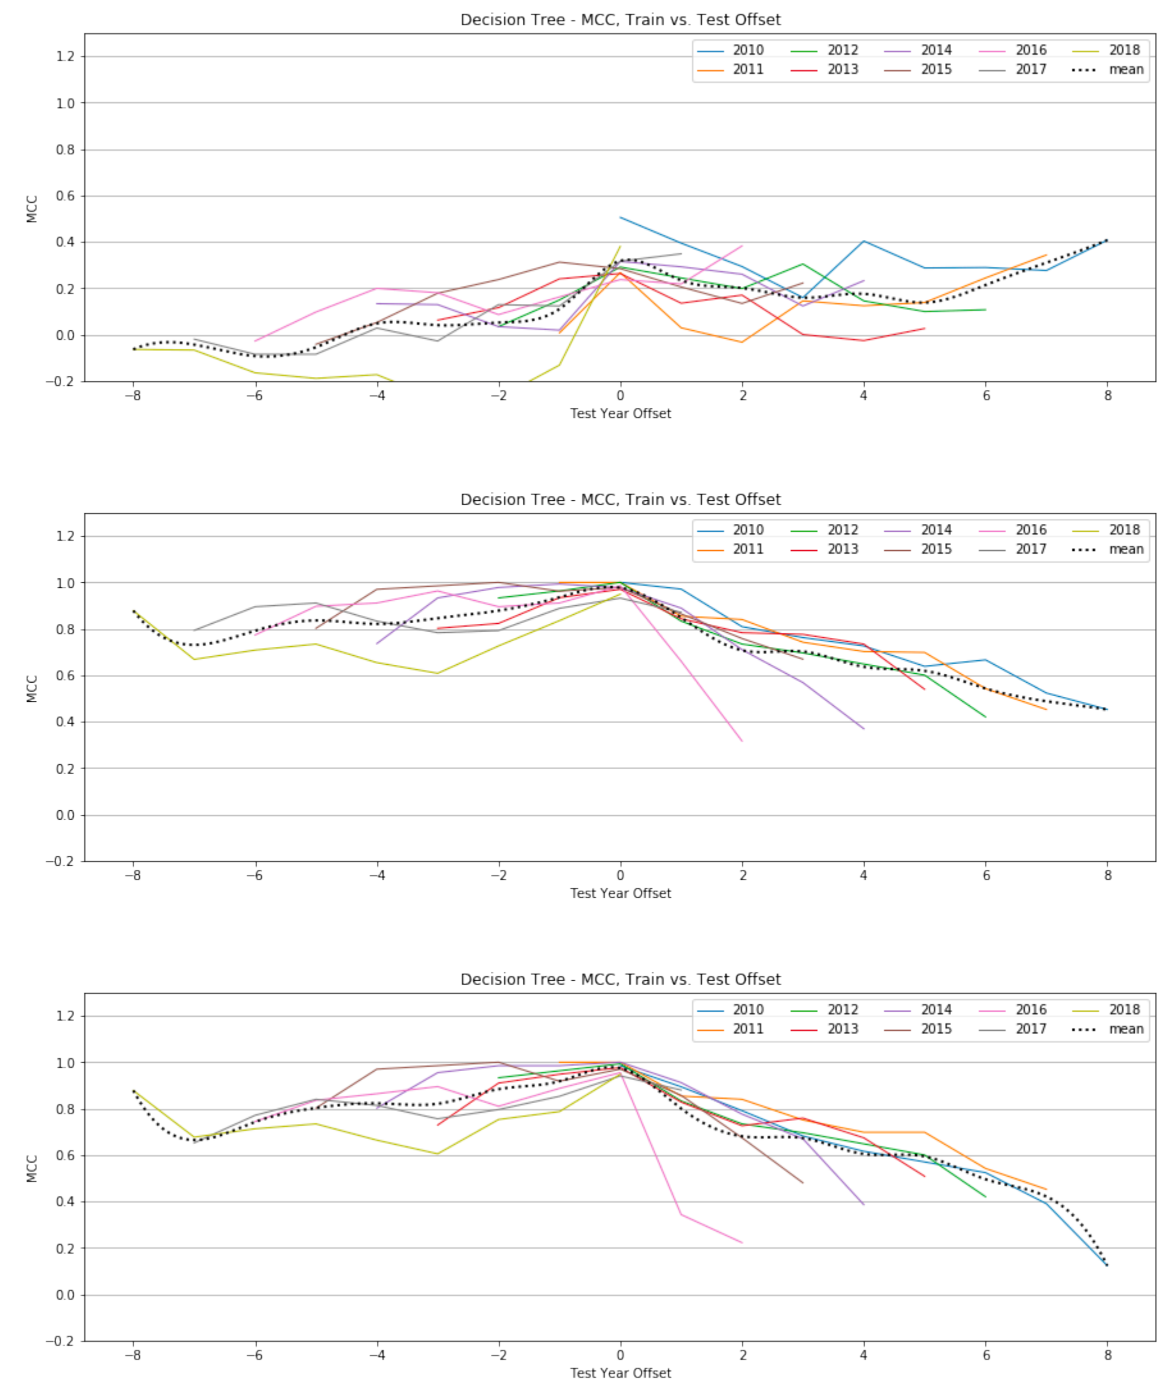

Figure 13: Decision Tree, MCC by year, Tweet-Only vs. User-Only vs. All Features

Random Forest

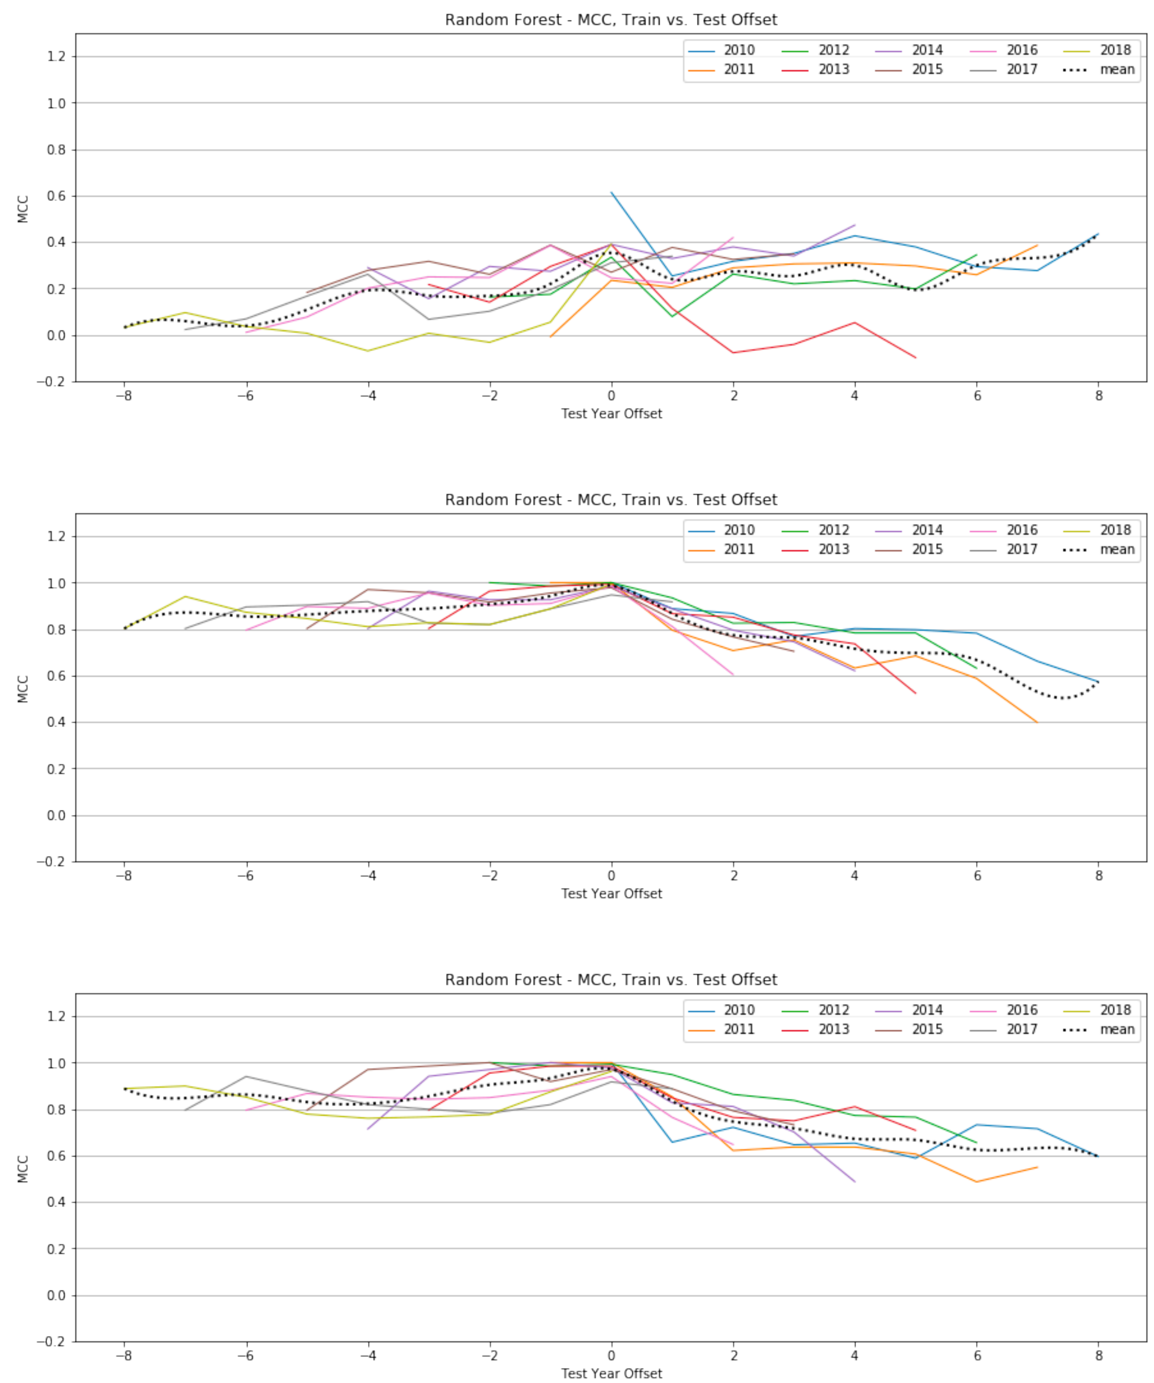

Figure 14: Random Forest, MCC by year, Tweet-Only vs. User-Only vs. All Features

Gradient Boosting

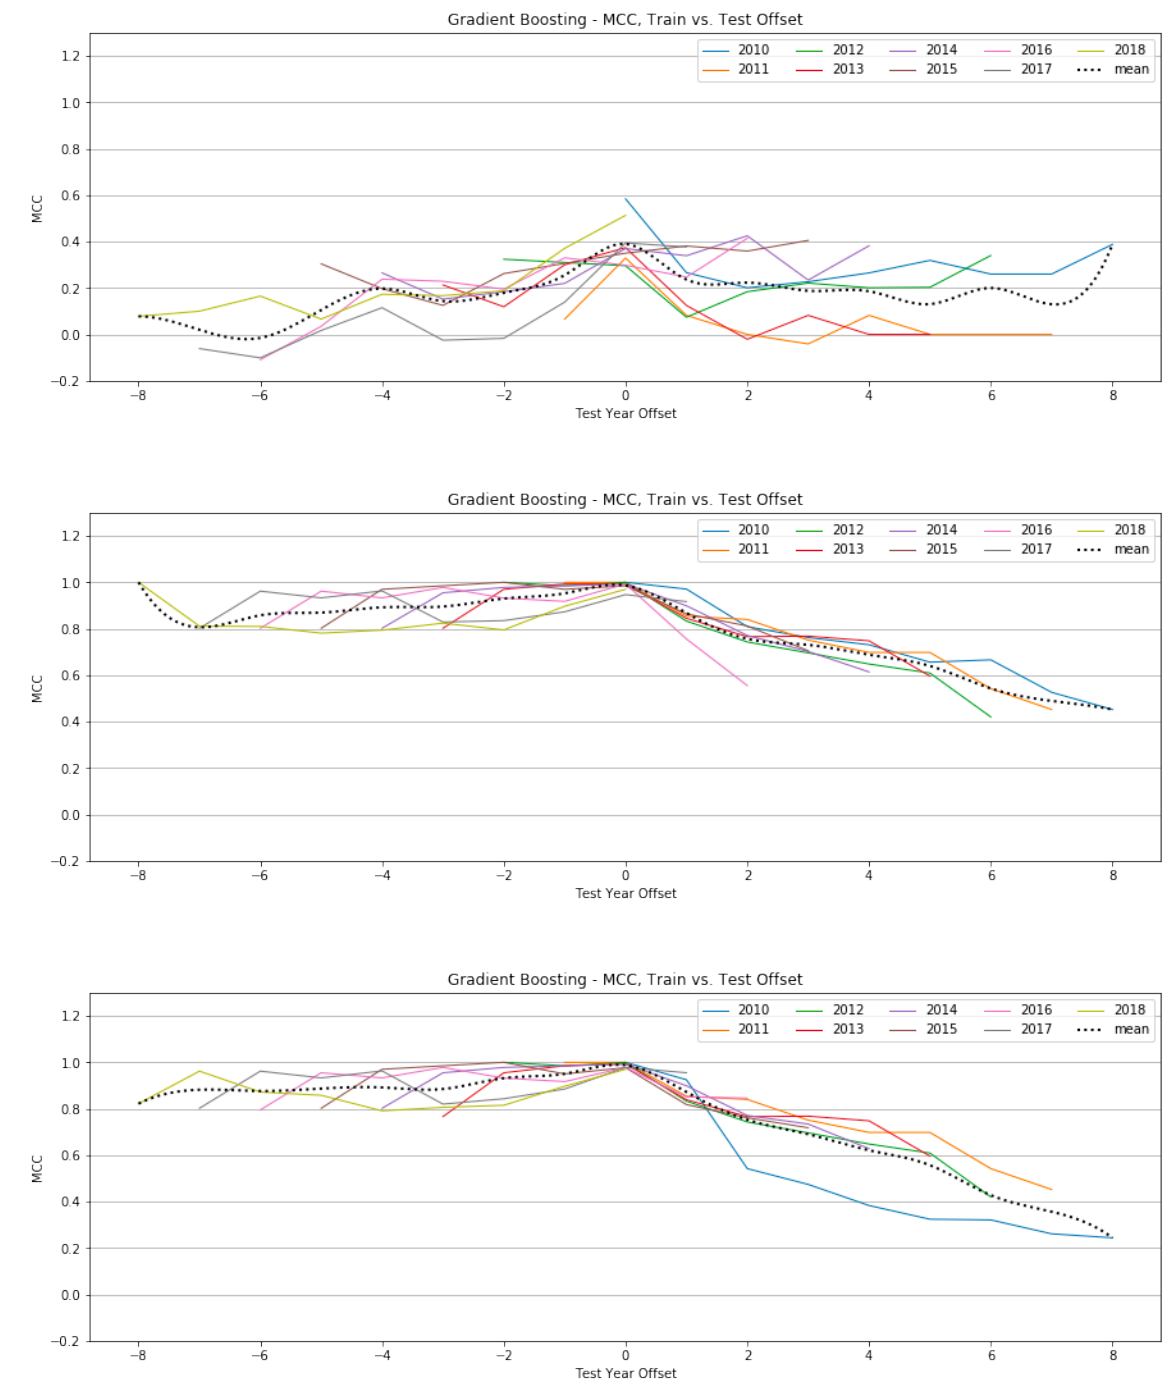

Figure 15: Gradient Boosting, MCC by year, Tweet-Only vs. User-Only vs. All Features

CAN THE "GORILLA" DELIVER?

Assessing the Security of Google's New "Thread" Internet of Things (IoT) Protocol

By: **Kenneth W. Strayer**, Deputy PM, Electronic Warfare and Cyber-PEO IEW&S US Army; Graduate, SANS Technology Institute, Graduate Certificate in Cybersecurity Engineering

SECURITY INCIDENTS ASSOCIATED WITH INTERNET OF THINGS (IOT) DEVICES HAVE RECENTLY GAINED HIGH VISIBILITY, SUCH AS THE MIRAI BOTNET THAT EXPLOITED VULNERABILITIES IN REMOTE CAMERAS AND HOME ROUTERS.

Currently, no industry standard exists to provide the right combination of security and ease-of-use in a low-power, low-bandwidth environment. In 2017, the Thread Group, Inc. released the new Thread networking protocol. Google's Nest Labs recently open-sourced their implementation of Thread in an attempt to become a market standard for the home automation environment. The Thread Group claims that Thread provides improved security for IoT devices. But in what way is this claim true, and how does Thread help address the most significant security risks associated with IoT devices? This article assesses the new IEEE 802.15.4 "Thread" protocol for IoT devices to determine its potential contributions in mitigating the OWASP Top 10 IoT Security Concerns. It provides developers and security professionals a better understanding of what risks Thread addresses and what challenges remain.

Photo Graphic Composite: Shelley Stottlar, Quanterion Solutions Inc., **Featuring Deposit Photos Stock Images:** by aa-w, Yur4you, and Finevector.

INTRODUCTION

Internet of Things (IoT) devices have become ubiquitous. The Gartner Group estimates that as of 2017, 8.4 billion connected IoT devices will be in use, not including smartphones, tablets, or computers, representing an increase of 30 percent from 2016 (Gartner, 2015). Security incidents associated with IoT devices have recently gained high visibility, such as the Mirai botnet that exploited vulnerabilities in remote cameras and home routers. Undoubtedly, the number of IoT devices will continue to expand, rapidly creating an ever-growing security concern.

The broad range of available protocols serves to compound the security problem associated with IoT. Developers have a choice among many competing technologies to include Wi-Fi, Bluetooth, ZigBee, cellular, Near Field Communication, Z-Wave and others, all of which come with inherent security advantages and disadvantages. In 2014, a consortium released a new networking protocol vying to become a market standard for the home automation environment. The new "Thread" protocol implements an IEEE 802.15.4 mesh network which is similar to ZigBee, but utilizes IPv6 technology with 6LoWPAN as its foundation. The developers advertise the standard as having "security and low-power features that make it better for connecting household devices than other technologies…" (Randewich, 2014). Google's Nest Labs recently open-sourced their implementation of Thread in an attempt to gain industry adoption as the standard for IoT.

The Thread Group makes strong claims for their new protocol. In one of their press releases, they claim that "Thread closes identified security holes found in other wireless protocols and provides worry-free operation" (Thread Group, Inc., 2014). But in what way is this claim true, and how does Thread help address the most significant security risks associated with IoT devices? Manufacturers have released very few IoT consumer products implementing the Thread protocol, leaving many unknowns. Most of the available information on Thread and the Google Nest implementation has been issued by the Thread Group and Google themselves, or are summations of the original developer's marketing material. Very little independent analysis has been conducted to assess the potential contributions of Thread to IoT security concerns.

This research paper will assess the new IEEE 802.15.4 "Thread" protocol for IoT devices to determine its potential contributions in mitigating IoT security concerns. It will evaluate these potential contributions utilizing the Open Web Application Security Project (OWASP) Top 10 IoT security concerns as a reference benchmark. The results of this study will serve developers and security professionals in better understanding what risks Thread may address and what challenges remain. It will help security professionals better analyze how devices are implementing the protocol at the data link, network, and transport levels.

Overview of the Thread Protocol

The Thread Group released the latest Thread 1.1.1 Specification on February 13, 2017. The specification provides extensive detail on the Thread protocol and claims to provide everything necessary to implement a Thread networking stack (Thread Group, Inc., 2017). The Thread protocol is described in the specification as "an open standard for reliable, cost-effective, low power, wireless device-to-device communications" (Thread Group, Inc., 2017, p. 1.3). The Thread standard is best referred to as a "network stack" in that it combines existing standards and protocols with specific implementation guidance to define the desired networking architecture.

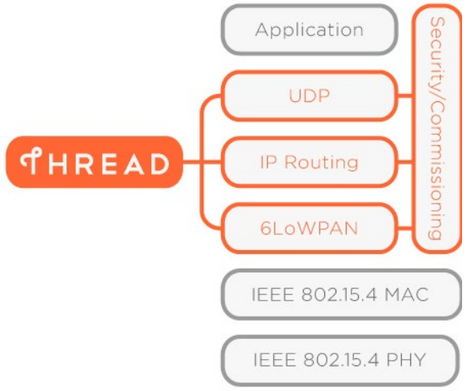

Figure 1: *The Thread specification defines existing protocols and standards for various layers of the interconnect model, with implementation guidance primarily focused on the network and transport layer.*

Various protocols were selected to meet the goals of Thread, to include support for IP-based addressing, use of existing hardware technology, scalability, low latency and power requirements, and simplified security (Thread Group, Inc., 2015b). As shown in Figure 1, the Thread networking stack primarily addresses the transport and network layers of the interconnect model, utilizing existing IEEE 802.15.4 radio components at the physical layer. Thread provides flexibility at the application layer, allowing a variety of market applications. According to the Thread technical overview, "Thread defines how data is sent in the network but not how to interpret it" (Thread Group, Inc., 2015c, p. n.p.). The application level flexibility becomes a critical point in assessing Thread's contribution in addressing IoT security concerns.

In addition to the physical radio, the Thread "stack" also specifies the 802.15.4 Media Access Control (MAC) layer for basic message handling and link layer control "for reliable messaging between adjacent devices" (Thread Group, Inc., 2015b, p. 4). The MAC layer also provides the primary encryption and integrity protection for Thread. IPv6-based addressing is fundamental to the Thread stack at the network layer (Thread Group, Inc., 2015b). To make Thread efficient over low power in a low-bandwidth environment, Thread utilizes "IPv6 Over Low Power Wireless Personal Area Networks" (6LoWPAN) for header compression and end-to-end fragmentation of messages. 6LoWPAN is an adaption layer that encapsulates the 802.15.4 packets for use over the IP network, providing a low overhead mechanism for forwarding multi-hop packets (Thread Group, Inc., 2015). Thread transport communications occur primarily via User Datagram Protocol (UDP) utilizing Datagram Transport Layer Security (DTLS) (Thread Group, Inc., 2015a). As described in the Thread 1.1.1 Specification, "DTLS is a variant of TLS with additional fields in the records to make it suitable for use over an unreliable datagram-based transport" (Thread Group, Inc., 2017, p. 1.4). Since Thread allows use of standard Internet Protocol, any number of additional IP-based services may also be leveraged.

As shown in Figure 2, the Thread network includes end-devices, routers, and commissioners. End devices can serve as routers that provide routing services as well as joining and security services for the network. They can also act as a border router that provides connectivity from the 802.15.4 network to external Wi-Fi or Ethernet (Thread Group, Inc., 2015b). A commissioner is the authentication server for the network, providing credentials for joining the network. Commissioners can be routers on the Thread network or external devices connected to the border router, such as a smartphone connected via Wi-Fi (Thread Group, Inc., 2015a).

Thread communicates as a mesh network to provide reliability, allowing individual devices to forward messages for other devices towards their end-destination. Each router device in the network has connectivity and current paths for all other routers in the network, communicating with Mesh Link Establishment (MLE) messages. The MLE messages establish and configure secure links, detect neighboring devices, and maintain routing costs between

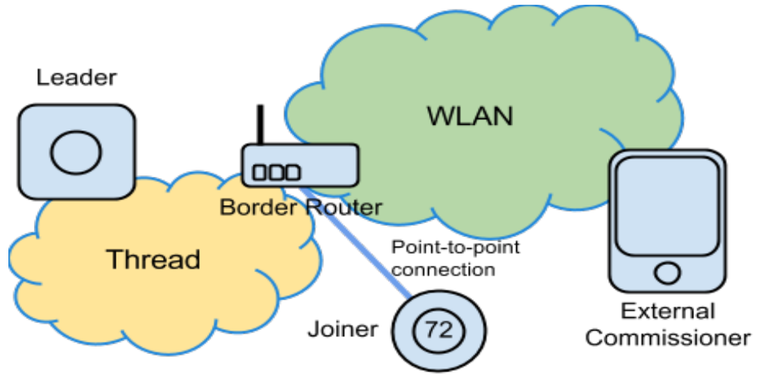

Figure 2: *One of several configurations that can be used with Thread for commissioning, authorizing, and joining of new devices.*

devices as the network changes (Thread Group, Inc., 2015b), creating an adaptable and secure transport architecture that requires no user interaction to maintain.

A device that seeks to join a Thread network must go through three phases: discovery, commissioning and attaching. A joining device discovers new networks through a beacon request. A responding beacon contains a payload with the network Service Set Identifier (SSID). Commissioning is typically performed by utilizing a commissioning application that is external to the 802.15.4 network, such as a smartphone connected via Wi-Fi through the Thread border-router, designed to force user-initiation for joining. Once the joining device has the required commissioning approved credentials and security material, it completes attaching to the Thread network via a Thread router (Thread Group, Inc., 2015b).

Overview of the OWASP Top 10 IoT Security Concerns

The Open Web Application Security Project (OWASP) is most commonly known for its Top 10 list of common web application vulnerabilities. But in June 2014, OWASP released its first list of Top 10 IoT security concerns. The foundation describes the OWASP IoT effort as being "designed to help manufacturers, developers, and consumers better understand the security issues associated with the Internet of Things, and to enable users in any context to make better security decisions when building, deploying, or assessing IoT technologies" (OWASP Foundation, 2017). The OWASP approach is to take a holistic approach to IoT security to include hardware interfaces, software configurations, network communications, and applications. As stated by one group of researchers, "Security is not an add-on feature; it must be built into the foundation of any given device. The level of security held by a device is derived from both the architecture and coding choices made by developers." (Sullivan & Sullivan, 2017, p. 14). While no single technology can be expected to resolve all the concerns across the various surfaces of an IoT device, the OWASP Top 10 serves as a useful framework to view the technology's contributions systematically and holistically.

RESEARCH APPROACH AND TEST BENCH

This study utilizes the OWASP IoT Testing Guide to develop an assessment and description of the Thread protocol's potential in mitigating each of the OWASP Top 10 IoT security concerns.

It is done primarily through analysis of available documentation to include the Thread Specification 1.1.1, various published white papers, and sample demonstration implementations provided by third-party vendors.

A hardware/software Thread test bench was built to include a control board, multiple radios, and a border router implementing the Thread protocol. This test bench was used to assess implementation of the Thread protocol to include live packet captures of component communications and the commissioning/association process. Analysis of the networking protocol in action provided opportunities to visually observe strengths and potential weaknesses as part of an end-to-end implementation.

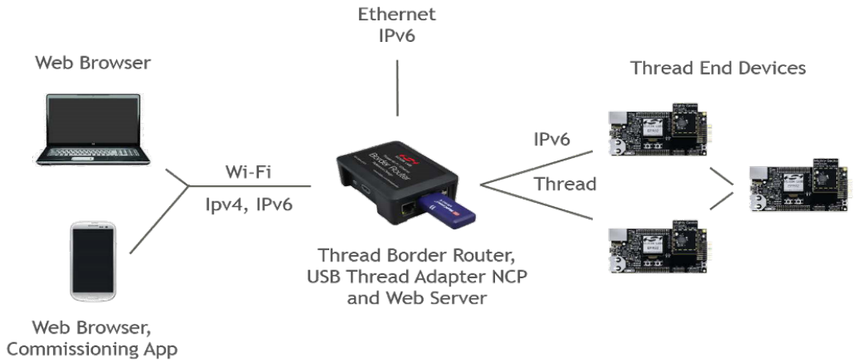

Figure 3: *The Thread Border Router System Components with EFR32 Mighty Gecko Wireless SoC Starter Kit served as the basis of the analyzed test bed.*

The SiliconLabs Mighty Gecko Wireless Starter Kit served as the basis for the test bench, along with the Thread software stack, sample code, and integrated debug adapter. As shown in Figure 3, multiple radio boards enable the creation of a demonstration mesh network utilizing a built-in IoT switch and light application to exercise the Thread software stack. The switch sends on/off and level control messages to the light in response to button pushes. A border router device allows management of the traffic between the Thread network and adjacent IP networks as well as the commission of the light and switch Thread nodes.

The SiliconLabs Simplicity Studio suite of tools provided sample applications, internal mesh network debugging, and packet tracing used to understand, analyze, and depict how the thread protocol implements security. Simplicity Studio's network analyzer is a graphical tool that displays network and node activity in real-time from either an untrusted perspective or with security keys to allow packet decryption and analysis. The network analyzer was used to validate the security assessment against the OWASP IoT concerns.

"Analysis of the networking protocol in action provided opportunities to visually observe strengths and potential weaknesses as part of an end-to-end implementation."

ASSESSMENT OF OWASP TOP 10 SECURITY CONCERNS

Insecure Web Interface

The first of the OWASP Top 10 IoT concerns deals with insecure web interfaces. As observed by the Infosec Institute, "The fact that your TV, toaster or baby monitor includes a web server is often a surprise" (Infosec Institute, 2014). Web interfaces are often poorly designed and insecure. Chances are, if a device has a web interface, it will also include default credentials. In 2012, the web application on TrendNet cameras was found to expose a full video feed to anyone who accessed it. While it included a secure sign-on capability, hackers quickly discovered that the authentication mechanism was just for show, and could easily be bypassed (Notopoulus, 2012). Besides default or weak passwords, one must assess the web interface for all the common vulnerabilities, to include cross-site scripting, SQL injection, lack of secure data transmission, and faulty account lockout mechanisms to prevent brute forcing (OWASP Foundation, 2016).

As shown previously, the Thread networking stack primarily addresses requirements at the transport and network layers of the interconnect model, providing broad flexibility at the application layer. The implication for Thread devices is that they are subject to all the same web interface vulnerabilities as any other IoT device. While Thread-based systems could have any number of web interfaces, the protocol defines two specific instances where a web interface is the standard implementation. First, the preferred method of commissioning new devices is through an external commissioner that allows a human administrator to manage joining to the Thread network. This device may be a smartphone or other device connected via a Wi-Fi network, or may be further extended to the cloud. Secondly, a Thread border router is typically employed to serve as a gateway between the 802.15.4 and the external Wi-Fi network. Both Thread features necessitate web interfaces for authentication and configuration, bringing all the common

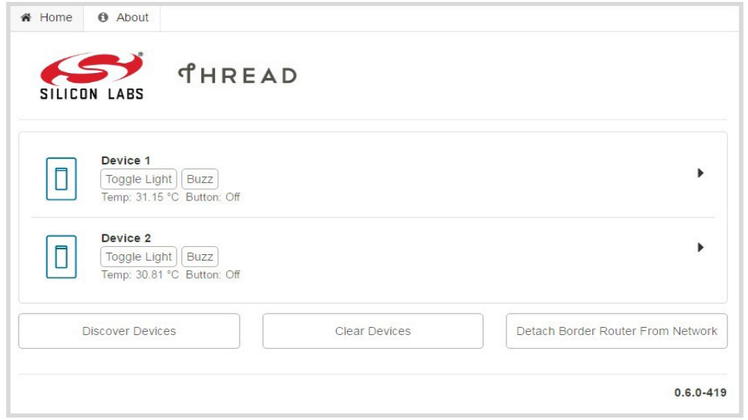

Figure 4: *The SiliconLabs Thread Border Router Sample Web Interface*

web vulnerabilities to the Thread IoT network. Before deployment, developers should test operational interfaces for account enumeration, weak or default credentials, account lockout, and fuzz test for SQL-Injection, cross-site scripting, or other flaws (OWASP Foundation, 2016).

The SiliconLabs test bench used in this study included sample applications and web interfaces for both the development border router (shown in Figure 4) and an Android-based commissioner. Considering the sample applications are for demonstration purposes only and not inherently controlled by the standard, this study did not conduct a complete vulnerability assessment. However, initial investigation of the Thread border router revealed a web interface with no credentialing interface, needing only a direct connection through its attached Wi-Fi access point. While the SSID of the access point was an arbitrary hex number, the passphrase was hard coded as "solutions" without any means to change the default setting. Additionally, as described in the quickstart documentation, when connecting to the border router, the commissioning application requires an admin password that is "set at compile time by the Border Router application and printed on stdout immediately after boot" (Silicon Laboratories, Inc., 2017b, p. 8). For the demonstration application, the developers set the admin password to "COMMPW1234."

For further analysis, the web interface in the border router was assessed using OWASP Zed Attack Proxy (ZAP), a free security tool used to scan web interfaces and applications for security vulnerabilities (OWASP Foundation, 2017a). The scan revealed two low-risk and one medium-risk alerts which included an issue with the X-Frame-Options header that could leave the web interface vulnerable to "ClickJacking" attacks. These problems would be considered critical flaws in any deployed product and demonstrate that the Thread protocol provides no significant contribution in addressing this particular OWASP Top 10 concern. In fact, Thread is agnostic at the application layer and is not designed to provide any web interface security enhancements.

"Default passwords provide easy access, while lack of mandatory password complexity can result in quick brute-force attacks."

Insufficient Authentication/Authorization

For IoT, authentication and authorization primarily involve weak or insufficiently protected passwords or credentials, or faulty authentication schemes. Default passwords provide easy access, while lack of mandatory password complexity can result in quick brute-force attacks. The Mirai botnet performs extensive scans of IP addresses to locate under-secured IoT devices with easily- guessable or default login credentials (Herzberg, Bekerman, & Zeifman, 2016). Some protocols, such as HTTP and FTP are notorious for passing credentials "in the clear" and can be easily sniffed and captured. These issues are all common in the implementation of IoT because developers often assume that interfaces will only be exposed on internal networks with minimal threat access (OWASP Foundation, 2017b). The OWASP security concern goes beyond credentials for web interfaces and addresses key management and network service authorizations. With poor key management or authentication, loss of a single node can compromise the entire system, or break the confidentiality and integrity of messages from other nodes (Sastry & Wagner, 2004).

Several credentials come into play in the ordinary operation of Thread devices, as well as in the joining and credentialing process. As discussed above, the web interface on commissioning devices, border routers, or the edge devices are not controlled by the Thread standard and may often be lacking appropriate security controls. However, the Thread standard does provide specific guidance on the implementation of transport and media access layer authentication and encryption. The standard claims that "Devices do not join the Thread Network unless authorized and all communications are encrypted and secure" (Thread Group, Inc., 2015b, p. 3). In order to achieve this, Thread utilizes a network-wide key at the Media Access Layer (MAC) to implement standard IEEE 802.15.4 authentication and encryption. The Thread standard describes the MAC layer encryption key as being "an elementary form of security

used to prevent casual eavesdropping and targeted disruption of the Thread Network from outsiders without knowledge of the network-wide key" (Thread Group, Inc., 2015a). However, the network-wide key is pre-shared and stored in non-volatile memory in the edge device. Any compromise of a Thread device could reveal the key and allow compromise of the network (Thread Group, Inc., 2015a). Also, distribution of the network-wide key to new devices on an IoT network is problematic. Asking consumers to enter authentication credentials into IoT devices that lack robust user interfaces adds complexity to the user experience, and the passing of credentials over unsecured connections would also be unacceptable. The Thread protocol commissioning process resolves these challenges.

During Thread network formation, the border router generates a random network master key. According to the Thread technical overview, the Thread software stack does not provide any mechanism for retrieving the key once created. If a Thread device is not yet a member of a Thread network and seeks to join, the thread protocol demands that the device first establish a secure Datagram Transport Layer Security (DTLS) connection with a Thread Border Router. Meanwhile, the commissioning device (an off-network smart phone, for example) establishes a secure DTLS session with the border router using a pre-determined commissioning passphrase. This passphrase is used to derive an enhanced key using key stretching (Thread Group, Inc., 2015a). A human operator then authenticates and authorizes the new joining device through the commissioning device. Once authorized, the border router provides the device the necessary security material to attach to the network over the secure DTLS connection that attackers cannot intercept. At no point does the commissioning device ever receive or hold the network security credentials, protecting from off-network exploitation (Silicon Laboratories, Inc., 2017a). Once joiner and border router exchange the network-wide key, the nodes utilize MLE messages "to establish and configure secure links, detect neighboring devices, and maintain routing costs between devices as the network changes" (Thread Group, Inc., 2015b).

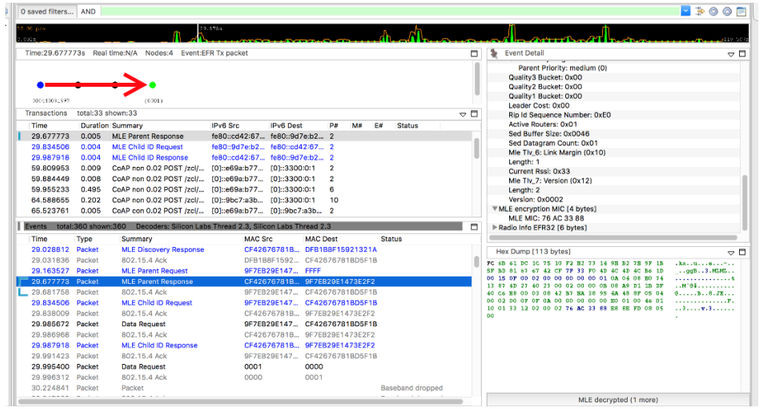

Figure 5: *The Thread commissioning process and network key exchange were observed in the study's test bench from a trusted perspective (with network keys to decode traffic) to confirm secure implementation.*

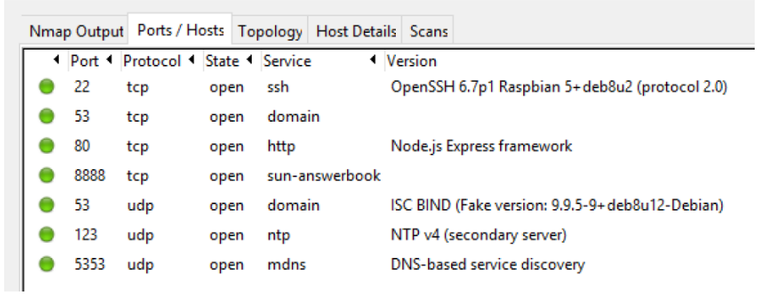

Figure 6: *A port scan of the test bed border router revealed several open ports without apparent functional requirements associated with the operation of the Thread system and may represent vulnerabilities.*

This paper's research included observation of the Thread commissioning process to confirm secure implementation. The test bed study included both trusted network captures (with internal network keys to decode traffic) and untrusted sniffing. Figure 5 illustrates the authentication and key exchange process with MLE Parent and Child requests and responses. Thread commissioning provides a secure means for distribution of key materials and simplicity in authorizing new devices to the network. The Thread border router commissioning process allows an autonomous self-configuring mesh protocol to implement MAC link-level security (Silicon Laboratories, Inc., 2017) in a simplified, user-friendly manner and significantly contributes in addressing the OWASP IoT concern for authentication and authorization.

Insecure Network Services

Weak network services in IoT devices can result in denial-of-service, or facilitate attacks on other devices. Devices may contain open ports that are unnecessary for their intended functionality. Developers often overlook these ports on IoT devices, assuming the network interfaces will not be exposed to external networks. Besides providing an access vector with weak credentials, these services can also often be exploited via buffer overflow or fuzzing (OWASP Foundation, 2017b). Attackers are very familiar with the vulnerabilities posed by insecure device ports and services. For example, the Mirai botnet went so far as to scan the infected device after initial infection and close off any open SSH, Telnet, or HTTP port services to prevent further infection from competing malware (Herzberg, Bekerman, & Zeifman, 2016). This is but one example of why insecure network services is a concern for IoT.

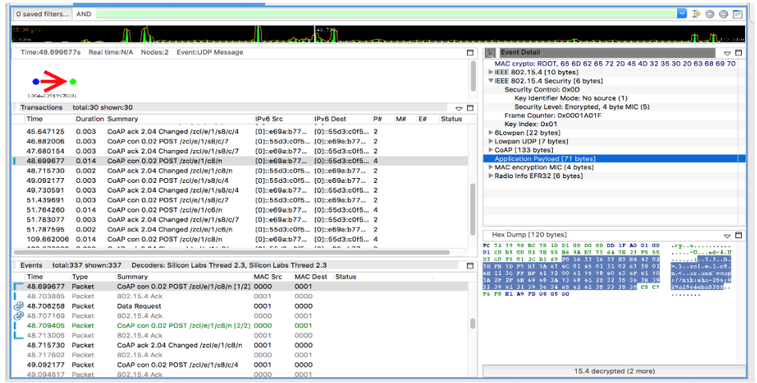

Figure 7: *When capturing data utilizing the pre-shared keys, you can decrypt the payload and read the 71-byte application payload containing the zcl 1/c6/n/ message to turn the light on or off.*

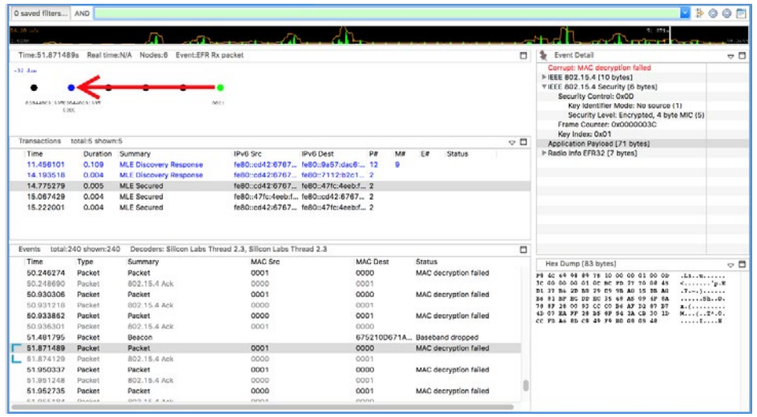

Figure 8: *Without access to the pre-shared key, the network sniffer cannot read the message payload protected by MLE message encryption.*

The Thread 1.1.1 Specification provides flexibility for implementation of various communication and commissioning topologies that may include border routers and off-network commissioning devices (Thread Group, Inc., 2017). Thread does not mandate specific hardware, software, or operating systems for such componentry, allowing configuration and deployment to support vendor-specific features while mandating consistency for the Thread specific functions (Thread Group, Inc., 2017). This flexibility poses immense challenges in securing network services and communication ports. The specification calls for various inter-device message exchanges utilizing UDP ports but has no limitation regarding the use of other UDP or TCP ports for functionality that is outside the constraints of Thread. The specification includes a series of SHOULD statements regarding firewalls and control of border router traffic but is focused on implementation of specific Thread processes broader security concerns.

The Thread test bed devices utilized in this analysis included a border router running Linux Raspbian Jessie Lite operating system as part of a standard Raspberry Pi computing device. A simple port scan of the device with ZenMap (Figure 6), revealed open ports without apparent functionality for the Thread network. While the border router must be able to assign IPv6 addresses to join edge-devices, it is not clear as to why a DNS server (TCP 53) is exposed on the interface facing the public internet or LAN router. TCP port 8888 and 5353 have no documented functionality for the test bed device. And while a Network Time Protocol service could be beneficial to an IoT network, exposure on the public side of the router only opens the system to potential additional exploits. For the test bench Linux configuration, IPTABLES was enabled providing a firewall service. However, the firewall was configured to allow all UDP traffic by default. As observed in this demonstration implementation, the Thread protocol does not provide significant contributions to address the OWASP IoT concern regarding insecure network services.

Lack of Transport Encryption

Transport encryption prevents data from being viewed as it travels across networks. Local networks are usually unencrypted and visible to anyone on the network. Wireless networks can often be misconfigured resulting in unauthorized access. IoT devices may utilize proprietary or weak encryption protocols. Lack of encryption can lead to exposure of data, but more importantly, it can provide critical information necessary to further compromise an IoT device or network (OWASP Foundation, 2017b). The use of encryption on IoT devices has been a constant challenge given the significant power drain associated with advanced features. To significantly contribute to

resolving this concern, an IoT standard must mandate accepted protocols that can operate in low-power environments.

Thread is advertised as a secure, power efficient standard for IoT. According to a Thread overview briefing, "Host devices can typically operate for several years on AA type batteries using suitable duty cycles" (Thread Fundamentals, 2017). To extend operations, Thread allows devices to sleep with adjacent nodes monitoring activities. The protocol mandates neighbor information exchange to include information on sleepy end devices and their sleep cycles (Thread Group, Inc., 2015b). These power management features allow the implementation of AES-128 link-layer security provided by the 802.15.4 MLE protocol. Additionally, since Thread utilizes 6LowPAN to encapsulate the 802.15.4 messages in IPv6, Thread allows the application to use any additional internet security protocol for end-to-end communication.

The study captured network traffic from the test bench and sample switch-light application to validate the operation of the Thread MLE message encryption. The initial data stream shown in Figure 7 was obtained by entering the pre-shared network key in the Simplicity Studio network analyzer module, allowing decryption of all message traffic to include the 71-byte application payload containing the 1/c6/n/ message to turn the light on or off. These captures demonstrate that if an attacker has access to the pre-shared keys (possibly through physically compromising and edge-node device), then the network confidentiality and integrity cannot be assured. When the same data stream is captured without access to the preshared keys, as shown in Figure 8, confidentiality and integrity is assured through MLE based encryption.

The Thread protocol implements both the 802.15.4 link-layer encryption as well as IP-based transport layer security enabled by 6LoWPAN and power management features to work in constrained environments. This

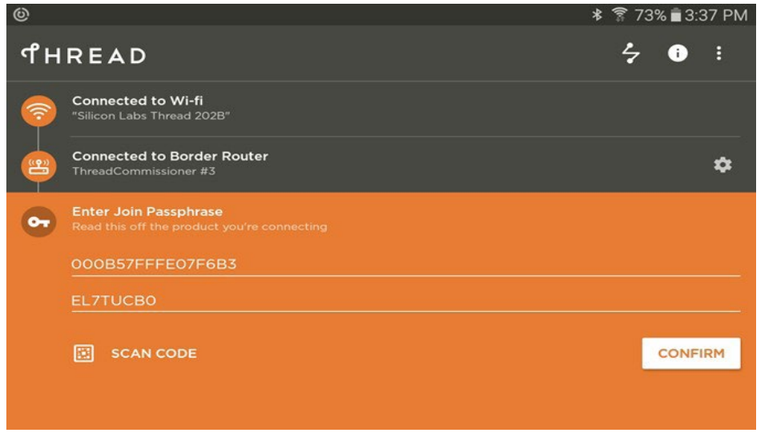

Figure 9: *The Android-based Thread commissioning app provided with the test bench system provided functionality for various consumer applications.*

combination provides a practical means to achieve a high-level of confidentiality and contributes significantly to addressing the OWASP security concern.

Privacy Concerns

Privacy concerns for IoT devices include both the collection and protection of personal data (OWASP Foundation, 2017b). Given the emerging, ubiquitous nature of IoT devices, personal data can go beyond financial and health records. IoT devices can provide insight into personal activities, preferences, and patterns allowing exploitation for nefarious purposes. Although the collection of personal data is an operational or functional concern, IoT privacy concerns magnify if a device has insufficient authentication, lack of transport encryption, or insecure storage of information (OWASP Foundation, 2017b).

In assessing privacy concerns for a Thread protocol device, security professionals must determine the amount of personal information collected, investigate the use of encryption at rest and in transit, and query end-user choices for data collection (OWASP Foundation, 2016). Apart from data transport, these items are application-level concerns that are not addressed by Thread. The simple switch and light application did not collect or store any private data. However, the included border router running a default version of Linux could be configured to log and store an unlimited collection of data. Except for Thread's default use of AES-128 encryption and its ability to leverage other secure transport protocols, Thread offers little in the way of contribution to the IoT privacy concern.

Insecure Cloud Interface

For most IoT devices, cloud-based data storage and access are integral to the required functionality. Off-premise storage of data leads to significant concerns for data protection. Insecure cloud interfaces often have weak credentials or allow account enumeration and manipulation of password reset mechanisms. The specific vulnerabilities are the same as the previous web interface concern which include default or weak passwords, lack of failed login lockouts, faulty password recovery mechanisms, or standard web-based vulnerabilities (OWASP Foundation, 2017b).

The Thread standard does not include any specific provisions for the implementation of cloud interfaces other than the ability to establish a commissioning device in the cloud and the inherent flexibility to implement other IP-based security applications and transport encryption. The Thread standard allows cloud-based data storage given the designed flexibility at the application-level, but the SiliconLabs test bench did not provide a specific cloud implementation. Giving credit for the Thread ease-of-use encryption and authentication mechanisms available for cloud interfaces, Thread only partially addresses the OWASP IoT security concern.

Insecure Mobile Interface

The presence of an IoT mobile interface implies remote access and potentially the control of the device over insecure public wireless networks. Developers of mobile devices are often pressured to simplify user access to the mobile interface given the constrained nature of the input mechanisms and screen size. Insecure mobile interfaces often have easily guessable credentials, lack two-factor authentication, and fail to encrypt passwords or other data during transport over public networks (OWASP Foundation, 2017b).

Most Thread systems would likely include a mobile interface for commissioning, and the potential for control and monitoring edge devices. The test bench for this study included an Android-based mobile application shown in Figure 9. The mobile interface included advanced features representative of several different consumer applications.

Once connected to the border router Wi-Fi network, the Thread mobile application searches for available Thread networks and requests the associated Thread administration password. As discussed in a previous section, the developers hard coded the Thread admin password for the test bench border router as "COMMPW1234." The sample application does not include two-factor authentication, or a means to change the admin password. However, data transport does take advantage of the previously described DTLS secure sessions mandated by the Thread specification providing full encryption during operation. While the Thread specification leaves much of the mobile interface design and security up to the developer, the communications mechanisms established for Thread commissioning simplifies security and partially addresses the OWASP IoT concern.

Insufficient Security Configurability

The ability to configure security options is essential in providing granular permissions for the access of data or controls for IoT devices. Broad access to certain data or functions on the IoT device may be a desirable feature for some applications, with the necessity of limiting access to administrative features such as the connection to new devices and password setting. To maintain high levels of security and privileged access, IoT devices require the ability to separate administrative users from ordinary users, and a means for monitoring and logging various security events (OWASP Foundation, 2017b).

The Thread specification provides little guidance on security configurations or separation of administrative and standard user features nor does it discuss monitoring or logging features at the border router. According the SANS Institute Internet of Things Survey, "system monitoring was cited as the second most common security control (65%) currently in use to secure Internet Things" (Pescatore, 2014, p. 19). However, system monitoring relies on the collection of central logs or host-based agents on edge devices. Thread does not control either of these capabilities.

"IoT devices require the ability to separate administrative users from ordinary users, and a means for monitoring and logging various security events."

The border router in our test bench is running a version of Linux operating system and is beyond the control of the Thread standard. Root access, storage encryption, communication ports, software updates, and logging are all independent variables that are addressed by the consumer application developer. Neither the web interface for the border router or the Android commissioning app had administrative controls, or a means to enable alerts or notifications. The OWASP IoT guidelines state the need for an active "security audit trail of mobile application interactions with the ecosystem" to include "robust logging and appropriate credentials to track interactions from mobile components" (OWASP Foundation, 2016a). Neither the assessed test bench or the Thread specification provide any capabilities to address this OWASP IoT security concern.

Insecure Software/Firmware

OWASP includes software and firmware security as a major IoT concern. According to OWASP, "the lack of ability for a device to be updated presents a security weakness on its own" (OWASP Foundation, 2017b, p. 31). First and foremost, devices must have mechanisms to allow easy updates as vulnerabilities are discovered and resolved. Additionally, software and firmware can be insecure if they contain hard-coded sensitive data or credentials. Depending on how systems distribute software and firmware updates, it is possible to intercept and compromise updates, unless mechanisms are in place to deny malicious software configurations, such as signing and verification of code (OWASP Foundation, 2017b).

The Thread specification includes standard message formats for reporting software versions for edge devices and border routers. However, it does not specify any means to manage or distribute software updates to these devices. The border router

in the test bench was running a version of Linux installed on an SD card. As part of the test bench analysis, the operating system was updated and patched. However, this process was problematic, requiring advanced knowledge including the ability to manually edit configuration files. Additionally, the analysis included the update of edge device firmware but required the use of a bootloader and flash program. More problematic, the test bench border router required a specific version of Linux, Raspbian Jessie Lite version 2016-05-31. Newer versions of the Raspbian operating system caused conflicts with the Thread border router services. The lack of a consumer-friendly means to update the software or firmware on these devices is indicative of a critical gap in the Thread protocol.

Additional investigation revealed hard-coded credentials in the edge-device application. The SiliconLabs documentation indicated that this was for ease of demonstration only stating, "While Thread applications deployed into the field are expected to use a randomly generated Master Key when starting the network, these Switch and Light examples applications use a hardcoded Master Key that can be found in the switch-implementation.c or light-implementation.c files…" (Silicon Laboratories, Inc., 2017c). While these samples are for demonstration purposes only, they indicate that IoT developers using the Thread specification are subject to the same problematic software or firmware security concerns.

Poor Physical Security

The last of the OWASP IoT Top 10 security concerns addresses poor physical security. If an attacker can easily disassemble a device or otherwise exploit the provided external ports, the installed operating system, and stored data become exposed. Attackers can modify devices for use in other purposes than those originally intended. One must review how easily device software can be accessed if any ports are present that are not necessary for normal operation, or if any administrative functions are limited or protected from physical tampering. Encryption of data at rest can further protect data on physically compromised IoT devices. (OWASP Foundation, 2017b).

"Encryption of data at rest can further protect data on physically compromised IoT devices."

According to the Thread 1.1.1 specification, "A Thread device MAY include multiple physical and media access control interfaces available for radio frequency or wired connectivity" (Thread Group, Inc., 2017, p. 3.13). The test bench border router included multiple ports to include Ethernet, USB, and a removable SD card for the operating system. The sample edge devices included USB and ethernet connections. A simple port scan of the edge devices revealed a TCP 4900 listening port, typically utilized for SQL client/services. In this case, the Ethernet port provides a means for debugging the radio application, but its presence represents an unknown security risk. Additionally, as detailed the Thread documentation, information is stored in non-volatile memory on the edge devices to facilitate rejoining to a network after a power loss or reset. The stored information includes the personal network identification, security material (each key used), and "addressing information from the network to form the devices IPv6 addresses" (Thread Group, Inc., 2015b, p. 19). Based on this analysis, the Thread standard is shown to have little contribution to addressing this OWASP IoT security concern, leaving secure design and testing in the hands of the consumer developer.

Assessment of Thread's Potential Contributions in Mitigating IoT Security Concerns

OWASP IoT Security Concern	Contribution
1. Insecure Web Interface	Minimal
2. Insufficient Authentication / Authorization	Significant
3. Insecure Network Services	Minimal
4. Lack of Transport Encryption	Significant
5. Privacy Concerns	Partial
6. Insecure Cloud Interface	Partial
7. Insecure Mobile Interface	Partial
8. Insufficient Security Configurability	Minimal
9. Insecure Software/Firmware	Minimal
10. Poor Physical Security	Minimal

Figure 10: *The Summary table depicts Thread's contribution in addressing the OWASP IoT Top 10 security concerns.*

CONCLUSION

This study assessed the new IEEE 802.15.4 "Thread" protocol for IoT devices to determine its potential contributions in mitigating IoT security concerns. Figure 10 provides the summary analysis of the protocol. Analyzing objectively, the Thread protocol provides significant contributions for authentication/authorization, as well as transport encryption.

The Thread standard was defined to provide an uncomplicated way to authorize and connect new devices while providing secure transport in a low-power environment. Implementation of 6LoWPAN provides options for implementation of additional IP-based security features. However, even though Thread represents an excellent standard to implement authorization and encryption on consumer IoT devices, its contribution to addressing the wider range of OWASP security concerns is limited. It only partially addresses privacy concerns, cloud, or mobile interfaces. The standard provides a minimal contribution to the remainder of the OWASP concerns, primarily due to Thread being a networking protocol abstracted from the application layer and the physical implementation/configuration of the IoT device. As shown in the study, device security is a factor of both the architecture and the implementation. This is analogous to the larger web application ecosystem that often leverages Transport Layer Security (TLS) using HTTPS. While TLS provides authentication and confidentiality, definitive security depends on the application itself. Thread has a role in providing a secure foundation for IoT systems, but it must be combined with well-conceived designs, thorough testing, and ongoing monitoring and patching.

REFERENCES

[1] Bush, S. (2015). Security is All in the Mesh Protocol. ***Electronics Weekly***, 15.

[2] Digi International. (2016, August 11). ***What Makes Thread a Secure Wireless Protocol?*** Retrieved from Digi: https://www.digi.com/blog/iot/what-makes-thread-a-secure-wireless-protocol/

[3] Gartner. (2015, November 10). ***Newsroom Press Release***. Retrieved August 05, 2017, from Gartner: http://www.gartner.com/newsroom/id/3165317

[4] Herzberg, B., Bekerman, D., & Zeifman, E. (2016, October 26). ***Breaking Down Mirai: An IoT DDoS Botnet Analysis***. Retrieved from Imperva Incapsula: https://www.incapsula.com/blog/malware-analysis-mirai-ddos-botnet.html

[5] Infosec Institute. (2014, November 10). ***How to Test the Security of IoT Smart Devices***. Retrieved from Infosec Resources: http://resources.infosecinstitute.com/test-security-iot-smart-devices/

[6] Krentz, K. F., Rafiee, H., & Meinel, C. (2013). 6LoWPAN Security: Adding Compromise Resilience ot the 802.15.4 Security Sublayer. ASPI'13.

[7] Notopoulus, K. (2012, February 3). ***Somebody's Watching: How a Simple Exploit Lets Strangers Tap into Private Security Cameras***. Retrieved from The Verge: https://www.theverge.com/2012/2/3/2767453/trendnet-ip-camera-exploit-4chan

[8] Olsson, J. (2014). 6LoWPAN ***Demystified***. Texas Instruments, Inc.

[9] OWASP Foundation. (2016, May 14). ***IoT Testing Guides***. Retrieved from OWASP: https://www.owasp.org/index.php/IoT_Testing_Guides

[10] OWASP Foundation. (2016a, May 14). ***IoT Framework Assessment***. Retrieved from OWASP: https://www.owasp.org/index.php/IoT_Framework_Assessment

[11] OWASP Foundation. (2017, August 12). ***OWASP Internet of Things Project***. Retrieved from Open Web Application Security Project: https://www.owasp.org/index.php/OWASP_Internet_of_Things_Project

[12] OWASP Foundation. (2017a, September 3). ***OWASP Zed Attack Proxy Project***. Retrieved from OWASP: https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project

[13] OWASP Foundation. (2017b, 08 21). ***OWASP IoT Top 10 PDF***. Retrieved from OWASP: https://www.owasp.org/images/7/71/Internet_of_Things_Top_Ten_2014-OWASP.pdf

[14] Palenchar, J. (2015). How Will Thread Stitch Together a Home Network? ***Twice.com***.

[15] Pescatore, J. (2014). ***Securing the "Internet of Things" Survey***. SANS Institute.

[16] Randewich, N. (2014, July 15). ***Google's Nest Launches Network Technology for Connected Home***. Retrieved August 5, 2017, from Reuters Technology News: http://www.reuters.com/article/us-google-nest-idUSKBN0FK0JX20140715

[17] Sastry, N., & Wagner, D. (2004). Security Considerations for IEEE 802.15.4 Networks. ***WiSE'04***.

[18] Silicon Laboratories, Inc. (2017, August 21). UG103.11: Application Development Fundamentals: Thread. Austin, TX.

[19] Silicon Laboratories, Inc. (2017a, August 13). UG116: ***Developing Custom Border Router***. Retrieved from SiLabs: https://www.silabs.com/documents/public/user-guides/ug116-border-router-ug.pdf

[20] Silicon Laboratories, Inc. (2017b, August 13). QSG102: ***Thread Border Router Add-On Kit Quick-Start Guide***. Retrieved from SiLabs: https://www.silabs.com/documents/public/quick-start-guides/qsg102-thread-border-router-kit.pdf

[21] Silicon Laboratories, Inc. (2017c, August 21). QSG113: Getting Started with Silicon Labs Thread. Austin, TX.

[22] Silicon Laboratories, Inc. (2017d, September 4). ***Thread Networking Solutions for the Connected Home***. Retrieved from https://www.silabs.com/products/wireless/mesh-networking/thread

[23] Sullivan, D., & Sullivan, J. (2017). The Internet of Everything is Everywhere; Testing Must Be Too. ***Information Security Magazine Insider Edition***, 13-16.

[24] Symantec. (2016, October 27). ***Mirai: what you need to know about the botnet behind recent major DDoS attacks***. Retrieved from Symantec Official Blog: https://www.symantec.com/connect/blogs/mirai-what-you-need-know-about-botnet-behind-recent-major-ddos-attacks

[25] Thread Group, Inc. (2014, July 15). ***Introducing Thread: A New Wireless Networking Protocol for the Home***. Retrieved from Thread Group: http://threadgroup.org/news-events/press-releases/ID/20/Introducing-Thread-A-New-Wireless-Networking-Protocol-for-the-Home

[26] Thread Group, Inc. (2015, July 13). Thread Usage of 6LoWPAN. San Ramon, CA.

[27] Thread Group, Inc. (2015a, July 13). Thread Commissioning. San Ramon, CA.

[28] Thread Group, Inc. (2015b, July). Thread Stack Fundamentals. San Ramon, CA.

[29] Thread Group, Inc. (2015c, October 4). Thread Technical Overview. Berlin.

[30] Thread Group, Inc. (2017, February 13). Thread 1.1.1 Specification. San Ramon, CA.

MR. KENNETH W. STRAYER has extensive industry and government leadership experience in development of Department of Defense tactical communication networks, software systems, and cyber warfare capabilities. After serving over 22 years as a military officer in the United States Army, Mr. Strayer led a large cyber and intelligence business development effort for the Science Applications International Corporation (SAIC). In 2017, he returned to government service in his current role where he is responsible for research, development, and sustainment of intelligence, electronic warfare, and cyber capabilities for the Army. In this capacity, he is responsible for a 275-person workforce and execution of over $400 million annually in programs. Mr. Strayer has a Bachelor of Science degree in Mechanical Engineering Technology from the University of Dayton and received a Master of Science degree from the Naval Postgraduate School in Monterey, California. He is a credentialed Project Management Professional (PMP) and holds a graduate certificate in Cybersecurity Engineering from the SANS Technology Institute.

REBOOTING LETTERS OF MARQUE

for Private Sector Active Cyber Defense

By: **Frank Colon,** Cyber Operations Attorney, 780th Military Intelligence Brigade (Cyber): Commander, U.S. Navy, Judge Advocate General Corps (ret.)

The views expressed in this paper are those of the author and do not reflect the official policy or position of the 780th Military Intelligence Brigade, U.S. Intelligence and Security Command, Department of the Army, Department of the Navy, Department of Defense, or the U.S. Government.

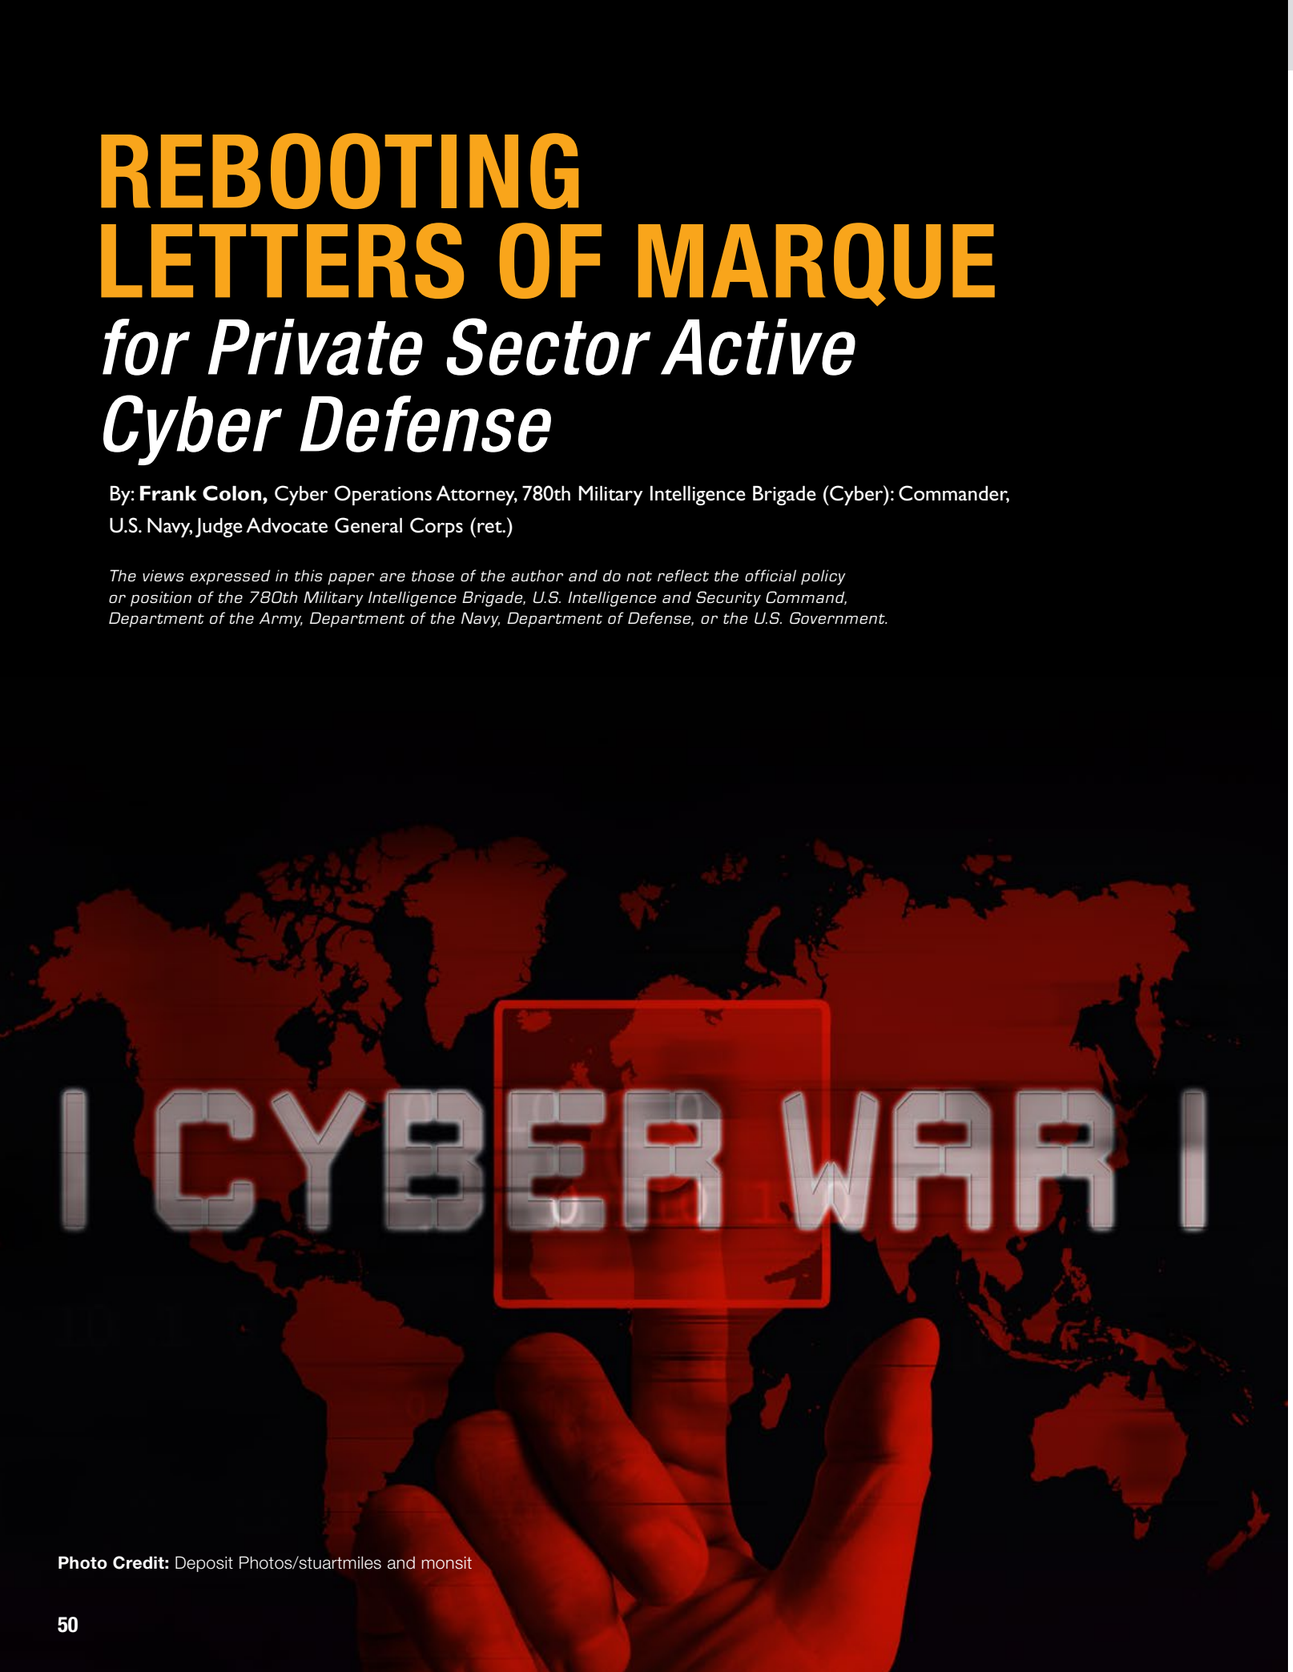

Photo Credit: Deposit Photos/stuartmiles and monsit

LETTERS OF MARQUE FOR PRIVATE SECTOR CYBER DEFENSE

Cyber assaults on U.S. Corporations will continue to increase until the United States articulates an enabling policy for the private sector to protect themselves by increasing costs to the hacker. The Center for Strategic and International Studies recently estimated cybercrime and espionage has caused $600 billion dollars' worth of damages.[1] U.S. Corporations are under cyber siege 24 hours a day in a "…borderless war that has impacted business across the world…."[2] On a daily basis hackers target businesses and individuals to steal data or damage digital systems. In many cases hostile foreign powers directly sponsor or otherwise enable the attackers. "In recent years, some foreign countries appear to have begun to operate in close cooperation with cyber criminals and the dividing line between where a criminal enterprise ends and where a nation state begins can often be difficult to determine."[3] Collectively these actors have virtually no consequence when attacking or attempting to attack private enterprise as all American private enterprise can do is lock the doors and hope for the best. Adding to the volume and complexity is the low cost of entry and lack of geographical boundaries.

Our adversaries, both nation state and criminal, have discovered that conducting offensive cyber operations against the United States in the "gray zone[1]" has incapacitated the United States. The gray zone creates an ambiguous security and legal environment. The gray zone permits nation states and their bad actor proxies to conduct unprecedented theft of intellectual property and personal information for illegal gain. While the reputation of the United States to defend itself from military aggressors is undeniable, we have yet to demonstrate our resolve and will to do so in cyber-space. As a result, nation states and criminals occupy the gray zone and dominate, we have; "…failed to keep pace with the threat."[4] Nation states in particular have used the gray zone to "… pursue their objectives while reducing the risk of triggering open warfare."[5]

In 2017, the average time for an intruder after entry to begin moving laterally to other systems in the network averaged 1 hour and 58 minutes.[6] Because of the speed at which a hack takes place, law enforcement cannot respond to an attack after it begins; only the victim has time to respond. Speed of relevance is critical to combating cyber-attacks. Businesses that come under cyber-attack have few legal or technical options beyond monitoring its network, fixing broken systems, and moving on. "…[U]nder U.S. domestic law, a private victim of a cyber-attack possesses a limited array of potential cyber responses. Digital self-defense, such as "hacking back," takes many forms from simply tracing an attack to identifying the culprit to damaging the hacker's machine. However, the same laws that prohibit hacking in the first place—such as the Computer Fraud and Abuse Act (enacted in 1986)—also prevent a company from striking back at maliciously motivated hackers."[7]

"On a daily basis hackers target businesses and individuals to steal data or damage digital systems. In many cases hostile foreign powers directly sponsor or otherwise enable the attackers."

A recent proposal to modify the criminal statute[2] only makes some aspects of hacking back a "defense" to criminal prosecution. A defense to prosecution does not prevent the matter going to a criminal trial. As a result, a hacking victim engaged in defensive actions could be prosecuted by cyber ignorant prosecutors and forced to hire uniquely qualified defense counsel at extraordinary costs. The proposed statutory modification also fails to address the potential for liability. Network and internet providers whose infrastructure were used to hack back might claim damages against the hacking victim who navigated those systems to engage in defensive actions. Since the proposed statutory modification does not mitigate serious risks, costly litigation, and tort liability to the private sector any participation is doubtful.

Similarly, U.S. Statutes providing for federal criminal charges for hacking are not effective against international hackers. A detailed search found a deficient number of foreign cyber prosecutions by the Department of Justice in 2015 through 2017. Moreover, no agency within the U.S. has principal responsibility for cyber security on behalf of U.S. Corporations. Recently, the Department of Homeland Security was given additional funding and authority to coordinate with local, state, tribal and territorial governments on security initiatives, while working to reduce and eliminate threats to critical infrastructure.[3] However, this new authority still leaves private sector not deemed critical infrastructure vulnerable. Finally, high costs and questionable effectiveness prohibit building of a cyber police force by the U.S. Government to protect the private sector. The Department of Defense provides cyber support to its industrial base under 32 C.F.R. 236. "The Pentagon reports more than 10 million efforts at intrusion each day."[8] In 2015 Senator Angus King complained during a Senate hearing: "We are in the cyber war with our hands tied behind our back. We would never build a destroyer without guns … you cannot defend, defend, defend, defend and never punch back."[9]

Imagine, one evening you are home with your loved ones and you hear your back door rattling. You go to investigate and you see an unknown "hacker" deliberately attempting to get in. You pick up your phone and call 911, but the operator tells you, sorry but we don't protect you from cyber intrusions. You hang up your phone and you hear a different noise at your window. When you investigate you

1 The gray zone has been defined as the space between peace and war.
2 Discussion Draft of the Active Cyber Defense Certainty Act 2.0: Section 1030 of title 18 United States Code.
3 Cybersecurity and Infrastructure Security Agency Act of 2018

see another hacker diligently working to gain entry. As you look to see the progress of the masked person at your back door, you hear a noise coming from your fireplace and it is not Christmas Eve. This silly hypothetical should give you a brief sensation of what it is like to own a cyber-network under persistent attack. If bad actors do steal valuable or sensitive data, victims will attempt to hold the custodian responsible under some tort theory of liability. If you believe this is the concern of some big corporation, remember your photos, medical, and other sensitive data is contained on those servers under unrelenting attack.

"Despite unrelenting attacks, endless data breaches, and data use abuses by social media and search engines, the internet has changed the world."

Despite unrelenting attacks, endless data breaches, and data use abuses by social media and search engines, the internet has changed the world. Web-connected devices provide access to instantaneous, unfiltered, global information. Even those in information hostile-nations, are supported by information freedom fighters who develop tools, tactics, and techniques to aid them in overcoming government restrictions on information. Not only has the internet contributed to the democratization of information, it also contributes to everyone's bottom line. It is estimated the internet contributed four trillion dollars to the world economy in 2016.[10] With informational and economic successes "…more than 20 billion devices are forecast to be connected, by 2020."[11]

THE INTERNET INFRASTRUCTURE

The internet infrastructure is comprised of multiple redundant interconnected digital networks owned by numerous companies and governments. In the United States, AT&T, CenturyLink, Cogent, Level 3, Sprint, and Verizon own the bulk of the U.S. internet infrastructure (backbone). These companies provide bulk service to Internet Service Providers (ISP) or to customers directly.

To assist with the discussions this article provides an oversimplified map of how the internet works. As previously stated, the internet consists of large infrastructure owners who deliver long haul digital routes for data flowing from and to different internet service providers who in turn deliver the data packets to end users. Because this process involves several hand-off points, no one internet operator or end user can see the entire transmission of data. Since the majority of data is legitimate, defenders have to be able to distinguish the bad data from the good in a never-ending stream of data that on its face appears legitimate.

Imagine looking at live images from a traffic cam of a particular stretch of road during rush hour traffic near a major city. There are thousands of cars (data) on the freeway (backbone infrastructure), some cars exit onto large local roads, (internet service providers (ISP), while some cars continue on the freeway out of the camera's view (data transfer to another backbone provider). In both cases, the backbone provider does not know what happens with the data once it exits its freeway or leaves its backbone boundary. That backbone provider knows the data came from X and went to Y. Neither points may be originations or final destinations making distinction, and, more fundamentally, tracking very difficult, all while happening at the speed of light. Returning to the traffic cam some cars that exited to large local roads controlled by ISPs now exit into large parking garages, (server farms or corporate networks), and finally some cars drive into private garages, (cyber citizens)[4]. In this case, the ISP does not know what happens to the data once it enters the server farm or cyber citizen device. The ISP knows the data came from X and went to Y. Unlike with someone watching the traffic cam, ISP providers may not readily be able to know where the data originated other than the immediate backbone provider it exited. The ISP would have to contact the backbone provider to see where that data entered that backbone. Similarly, the large corporate end users or server farms can't see past its ISP. Despite the interconnected aspect of the internet the hand-offs create knowledge gaps that bad actors exploit, and there is no corresponding mechanism or statute currently in force to attempt to mitigate or deter bad actors.

LETTER OF MARQUE

United States Constitution: Article I, Section 8, Clause 11 in the United States Constitution states: "The Congress shall have Power ... To …, grant Letters of Marque and Reprisal, and make Rules concerning Captures on Land and Water;"[12]

A Letter of Marque is a government license authorizing a private person or entity to take an action on behalf of the issuing government, which could include permission to cross an international border, and in some cases after review by a court transfer title of the goods captured to the license holder as a "prize."

Historically, to request a Letter of Marque, a ship-owner would apply stating the name, description, tonnage, and force (armaments) of the vessel, the name and residence of the owner, and the intended number of crew, and tendered a bond promising strict observance of the country's laws, treaties, and of international laws and customs. The commission was granted to the vessel, not to its captain, often for a limited time or specified area, and stated the enemy upon whom attacks were permitted. For instance, during the Second Barbary War President James Madison authorized a brig[5] named the *Grand Turk* to cruise against Algerian vessels, "…public or private, goods and effects, of or belonging to the Dey[6] of Algiers".[13] The East India Company (a British Company) arranged for letters of marque so that, should they have the opportunity to take a prize, they could do so without being guilty of piracy. However, the United States has not issued a letter of marque since the War of 1812. Interestingly, in December 1941 until 1942, Goodyear's commercial L class blimp *Resolute* operating out of Moffett Field in Sunnyvale, California, flew anti-submarine patrols. As the civilian crew was armed with a rifle, many believed this made the ship a privateer, and that she and sister commercial blimps were operated under letter of marque[7] until the U.S. Navy took over this patrol.[14]

CYBER LETTER OF MARQUE

A 21st century Cyber Letter of Marque would not grant U.S. Corporations (private entities whether public or privately held) the right to capture a prize. However, a Cyber Letter of Marque would permit the right of self-defense outside of a corporation's network borders. Currently, U.S. Corporations protect and defend their network only after penetration by the bad actor – not the preferred position for defense. The strategic advantage and likelihood for success has clearly passed to the bad actor. A Cyber Letter of Marque would permit (vetted, trained, and bonded) American businesses to watch outside its network to look for pre-attack indicators and when attacked respond beyond the network borders. A cyber letter of marque provides a mechanism to facilitate a more robust and effective cyber defense for U.S. Corporations.

Given the inherent complexity of detecting nefarious cyber activities, no one specific level of internet provider or corporate user can singlehandedly deploy a Cyber Letter of Marque. Combating cyber threats requires custom tailored Cyber Letters of Marque with applicable authorities specific to the unique response possibilities for each entity involved in the transfer of internet data to respond and repel attacks and determine origin. Cyber Letter of Marque authorities would be tailored to each recipient and only after careful consideration of the strategic consequences and capabilities of the Cyber Letter of Marque holder. Robust communications between the layers of internet operators, and corporate end users, with unique authorities at each layer create opportunities for collective cyber response actions and reducing the unchallenged volume of cyber intrusions and attempted intrusions before an attack gains momentum.

"A Cyber Letter of Marque would permit (vetted, trained, and bonded) American businesses to watch outside its network to look for pre-attack indicators and when attacked respond beyond the network borders."

However, just giving U.S. Corporations expanded authorities will not solve the relentless volume of hackers. Nor can a sole government solution protect everyone on the internet. A whole of Nation solution is required. Department of Homeland Security (DHS): National Cyber Security and Communications Integration Center (NCCIC), "[s]trives for a safer, strong Internet for all Americans by responding to major incidents, analyzing threats, and exchanging critical cybersecurity information with trusted partners around the world."[15] However, "major incidents"[8] are not the norm. Gray zone cyber incidents are the norm and providing measurable success for our adversaries and billions of dollars in domestic damages. DHS announced a new center to be known as the National Risk Management Center and will provide a centralized home where firms (likely critical infrastructure) can turn for cybersecurity solutions.[16] A cyber drill dubbed "Jack Voltaic 2.0" was conducted in Houston, Texas in July 2018[17]. The exercise demonstrated gaps in operational and legal authorities. The Chief Technology Officer at the Houston Police Department said: "The assumption is that [the Department of Homeland Security] will be there, but that's not entirely the case."[18] Readiness teams sent by the DHS National Cybersecurity and National Communications Integration Center, "…can give advice, but not a lot." Bell said.[19] If Cyber Letter of Marque holders are to be truly successful and change the paradigm of the gray zone, both private sector and government need to establish persistent/enduring approach to countering gray zone cyber incidents. By expanding the National Risk Management Center to support Cyber Letter of Marque holders in a Cyber Fusion Center, connects stakeholders in real time facilitating synchronization of efforts and effects.

Liaisons: Private sector participants issued a Cyber Letter of Marque will assign cleared representatives who will be physically located in a Cyber Fusion Center as liaisons who have instantaneous reach back with the Corporate Network Operations Team. Similarly, Tier Three (discussed below) Federal Agencies and other relevant Federal and State Agencies will also have liaisons with reach back capabilities to Law Enforcement and U.S. Military cyber operators and other government resources. Co-locating representatives permits real time connection to the whole of Government with the private sector. Now when gray zone cyber incidents are initially detected, all relevant parties are seeing the response action in real time. Participating members of the Intelligence Community and Law Enforcement will inject relevant information that could facilitate the response action. If the attack is multi-pronged, other relevant agencies and private sector participants response time is significantly reduced. The whole of government and private sector cannot

4 Cyber Citizen: A human, regardless of location or jurisdiction, who uses technology in an appropriate and lawful manner.
5 A brig is a sailing vessel with two square-rigged masts.
6 **Dey** was the title given to the rulers of the Regency of **Algiers** (Algeria) and Tripoli under the Ottoman Empire from 1671 onwards.
7 No Letter of Marque was issued to Goodyear.
8 DHS has not defined "major incident."

effectively work together after a cyber incident has started. However, if private sector and relevant government agencies work together in a fusion center 24/7/365 and participate in joint exercises to test workflows, then a truly efficient nexus can be created to rebuff gray zone attacks.

"Congress holds the power to issue Letters of Marque under the United States Constitution. That authority could be delegate to the Department of Commerce or other appropriate agency to issue Cyber Letters of Marque."

INTERNATIONAL TREATY

Paris Declaration: In 1856, Britain, France, and other world powers met in France to discuss concerns arising from wartime maritime law.[20] In response to the United States' and others' effective use of privateers the Paris Declaration of 1856 was a document attempting to ban privateering. However, the United States refused to sign the agreement.[21] The Paris Declaration states that it is not a universal ban on privateering and only applicable to signatory nations at *war* with *other signatory nations*, [emphasis added] and does not have the authority to police the actions of non-signatories.[22] Accordingly, the plain language of the document does not apply to the United States. "Additionally, the Declaration clearly pertains and limits itself to maritime law. Since a cyber letter of marque regime is not grounded in maritime law and letters of marque are specifically authorized in the United States Constitution, it is permissible under international law, Paris Declaration notwithstanding, to issue cyber letters of marque."[23]

DEPLOYMENT OF CYBER LETTERS OF MARQUE

Congress holds the power to issue Letters of Marque under the United States Constitution. That authority could be delegated to the Department of Commerce or other appropriate agency to issue Cyber Letters of Marque. Prior to issuing a Letter of Marque, a Cyber Letter of Marque Program would be developed for U.S. Corporations to apply and participate. The enrollment would be voluntary and participation costs borne by the U.S Corporation. After successfully completing training, the private sector employees would receive certification by U.S. Cyber Command and Federal Law enforcement. Upon certification corporate employees will participate in (sector specific) exercises that require the skills they have learned to be deployed in a safe training environment. After completion of the program a Letter of Marque for Active Cyber Defense would be issued to the U.S. Corporation. The Letter of Marque would detail specific authorities, and any limitations.

Participating private sector employees would be in two tiers: One tier using unclassified tools and techniques, and tier two using higher level cyber tools requiring a Department of Defense security clearance.

TIER 1. PRIVATE SECTOR RESPONSE

Pre-Approved Unclassified Tools: As an example, for the financial sector, U.S. Cyber Command, Department of Homeland Security, along with Treasury will establish a set of "response actions"[9] that are exempt from U.S. laws that prohibit "hacking back." Pre-approved response actions will not be classified, reducing the number of employees who require a security clearance, and maximizing the number of certified corporate network responders. Companies can develop proprietary responses that can be cleared of criminality[10] in advance and could also be licensed, or shared for a fee or free. Companies that hire, train, and retain the best responders can market their enhanced security, or recover development costs under license or fee arrangements. The Federal Government for its part will always have a no fee license for defense of essential Federal systems. While proper network configurations, and good network and system hygiene, create an environment to repel high volume low threat cyber-effects, the ability to respond directly to illegal hacking will alter the cost benefit calculation for the hackers.

No U.S. Criminal Liability if using approved tools and techniques Foreign Criminal Liability would require a nation to acknowledge that a hacker operating within its geographical borders was victimized by the U.S. Corporate response. While legally possible, it places the charging nation in an embarrassing international position of raising criminal charges based on the claims of a criminal or state hacker against an actual victim.

International Law: Applicability of the Paris Declaration requires the use of a letter of marque for *maritime* purposes, *between a signatory nation* against *another signatory nation* who are *at war.* Physical presence or the conduct of business by a U.S. Corporation in a signatory nation is not enough to activate the Paris Declaration. However, U.S. Corporations would not be permitted to launch cyber defense actions authorized by a U.S. issued letter of marque in other nations without host nation consent. Civil liability remains to ensure private sector participants hire, train, and supervise skilled Tier One responders. Cyber Letter of Marque holders will carry a Bond to cover any civil liabilities or damages.

TIER 2. CLEARED PRIVATE SECTOR EMPLOYEES RESPONSE

Pre-Approved National Level Classified Tools: Similar to Tier One, relevant Federal Agencies will pre-approve and assign to specific private sector participants classified tools that can only be used by cleared private sector personnel when an attack escalates beyond Tier One.

Deny Internet Access to Infected Devices: Devices with internet access within the U.S. are either willingly or unwittingly participating in the cyber-effect. Tier Two responders can temporarily deny internet access to those infected or participating devices in order to contain the attack. Internet access blocking is only authorized when necessary to restore network functionality or to aid in the pursuit of the bad actor. Internet access blocking is not authorized for any compromised Federal systems, hospitals, or critical infrastructure. Internet Access blocking is permitted against privately owned computers with compromised systems and active attack or effect participation. Internet access denial is only authorized to permit enough time for the targeted network to be restored or 24 hours, whichever is less. If more than 24 hours is needed or the targeted network is critical the Federal Government, (law enforcement or DoD) will assume the active defense under Tier Three.

Since Cyber Letter of Marque authorities are tailored specifically for each participant, when they work together the specific authorities provided to each can enhance the overall response action when coordinated. Therefore, a financial institution that is responding to an attack under its Cyber Letter of Marque authorities may need the assistance of one or more network providers to coordinate the response.

If cyber defense conducted under a Cyber Letter of Marque begins to expand to a sensitive nation, or sensitive target, a decision by a government representative at a Cyber Fusion Center will be made in real time. A senior watch officer at a Cyber Fusion Center will determine which federal agency assumes the response action. Once assigned in real time that Federal Agency will follow the Agencies' existing command and control authorities. If the private entity is going to pass the response to the Federal Government then the hand off will occur outside the private sector network boundary. Keeping the Federal Government outside the private sector network eliminates the potential for the Federal Government to cause damage to private sector systems and protects the privacy of the private sector clients and data. When the Federal Government assumes the response position we enter the third tier listed below.

No U.S. Criminal Liability if using approved advanced tools and techniques. Foreign Criminal Liability would require a nation to acknowledge that a criminal or state hacker operating within its geographical borders was victimized by the U.S. Corporate response.

International Law: Some nations might claim U.S. Corporations are acting as cyber mercenaries as they are now using State level tools and techniques. The most widely accepted definition of mercenaries is found in Article 47(2) of Additional Protocol I of the Geneva Conventions, It sets forth the conditions that must be met:

> Special recruitment to fight in an armed conflict,
>
> Directly participates in hostilities,
>
> Is motivated by private gain, and is promised by a party to the conflict of material compensation in excess of that paid to combatants of similar ranks and functions,
>
> Is neither a resident nor national of a party to the conflict,
>
> Not a member of the armed forces who are involved in the conflict,
>
> Not sent by another state of official duty as a member of its armed forces.

In short, no. Applying Additional Protocol I, several conditions are not met in order to declare Cyber Letter of Marque holder mercenaries. Cyber Letter of Marque holders are protecting their own private property, even if the company is publicly traded. If successful they do not gain anything more than restored dominion over that which they already own. The Federal Government does not pay Cyber Letter of Marque holders to participate even if they are successful. In fact, Cyber Letter of Marque holders pay to participate and for the training of their personnel. Cyber Letter of Marque authorities are only available to U.S. Corporations, who are in fact residents in the nation. Finally, DoD contractors using Cyber Letters of Marque are not members of the armed forces, nor sent to conduct offensive operations, but are conducting defensive actions. Accordingly, Cyber Letter of Marque holders to include DoD contractors are not mercenaries under international law regardless if participating in a declared armed conflict or not.

TIER 3. FEDERAL LAW ENFORCEMENT / DEPARTMENT OF DEFENSE CYBER RESPONSE ACTIONS

Federal law enforcement and DoD Cyber Forces who have been following the cyber engagement can make recommendations to the private sector team, or take the response over deploying advanced Nation state level tools, effects, and techniques. Since both DoD and Federal Law enforcement have been involved from the beginning it is easy to determine which agency has primacy over the cyber response. If the cyber effect originated from the U.S. or friendly western nation, and after the attack is repelled, federal law enforcement will organize the evidence already collected from the engagement and proceed as a criminal case. If the cyber effect originated from an adversary or unfriendly nation, the DoD will have primacy over the event and respond accordingly. In this model real time Federal monitoring expedites the "law enforcement / military" decision point. Most importantly, the cyber effect has been rebuffed and only when the private sector was overwhelmed or the response actions required are outside the scope of the Cyber Letter of Marque, will the Federal Government respond. As a result of this model corporations bear the costs of Tier One and Two responses and only when an active cyber defense is transferred in Tier Three does the federal

9 Response actions: A menu of pre-authorized options for a particular industry sector that would be outlined in a proprietary annex to the Cyber Letter of Marque.
10 A review, testing, and vetting process would need to be created to support Cyber Letter of Marque holders.

government participate actively. The longer private vetted U.S. Corporations pursue bad actors the greater the likelihood they will succeed in determining the origin of the bad actor and repel the attacks.

attackers is inadequate..." Therefore, if active cyber defense under Cyber Letter of Marque authorities focuses on increasing costs and reducing gains to the hacker by impacting time and effort of the hacker, attribution is less relevant.

Attribution: Deputy Secretary of Defense William Lynn wrote in 2010, "Whereas a missile comes with a return address, a computer virus generally does not."

No criminal or civil liability for the private sector participants as they are out of the fight. No change to existing Federal Tort Law.

EFFECTS OF A PRIVATE SECTOR CYBER LETTER OF MARQUE

Attribution: Deputy Secretary of Defense William Lynn wrote in 2010, "Whereas a missile comes with a return address, a computer virus generally does not."[24] Attribution as it relates to cyber-attacks is an epic point of frustration. Alexander Melnitzky argued that attribution may be a bit overblown in his article "Defending American Against Chinese Cyber Espionage Through the Use of Active Defenses."[25] Others argue that without attribution your right of response is limited. This belief is based on the punishment aspect of deterrence. One must know who is attacking in order to deliver an appropriate measure of justice in response to the actor. "Classical deterrence theory rested primarily on two main mechanisms: a credible threat of punishment for an action; and denial of gains from an action."[26] As demonstrated no credible threat of punishment (i.e. incarceration) exists for international hackers. Accordingly, a better approach to deter hackers is to focus on the cost-gain analysis. "Deterrence is a function of the total cost-gain expectations of the party to be deterred, and these may be affected by factors other than the apparent capability and intention of the deterrer [sic] to apply punishments or confer rewards."[27] "[T]his means that a defensive effort is inadequate for better cybersecurity a strategy that does not impose consequences on

Cyber Citizens: Historically Letters of Marque have been used against governments, corporations, pirates11, and private individuals of other nations. However, what happens when attacks appear to originate from the United States. Botnet attacks allow remote control of computers whose owners have not properly protected and updated their connected device or even aware of the improper use of their device. As a result, an attack or effect might initially look like it is from computers within the United States. This fact along with existing federal criminal law has directly impeded U.S. Corporations' ability to actively defend networks. Cyber Letter of Marque holders, specifically internet service providers, will be permitted to temporarily deny internet access to non-federal compromised systems for 24 hours or less, only for the time required to contain the attack. While on its face the temporary loss of access to the internet by cyber citizens and private corporations sounds like a bad unintended effect, the never-ending cyber-attacks of the 21st century also must be addressed. Since internet service is a commercial product the temporary loss of access is reasonable when a poorly maintained computer system is part of an attack that denies access to thousands of innocent cyber citizens banking sites.

INTERNATIONAL REACTIONS

While I found no violation of international law in resurrecting Letter of Marque for cyber, that does not guarantee an absence of international reactions. As has been demonstrated recently cognitive warfare is alive and well. Cognitive warfare is about controlling the decision cycle. Those who fear active cyber defense and fear delegation of authority will conjure images of global escalations arising from corporate cyber defense measures. Others who want to keep the United States from acting will claim we are increasing hostilities in cyberspace. "Concern about escalation should not lead to timidity or indecision. This is a contest of wills and our opponents will use threats to bluff us into continued inaction. However, the same political constraints on the conduct of warfare that hamper the U.S. ability to respond to opponent cyber actions using military [kinetic] forces will also hamper them. For a better defense, the U.S. will need to become more comfortable operating in the "gray zone" that our opponents now inhabit." We have been victims for far too long. We developed the technology that underpins the global internet and because of a complete lack of will, the world has surpassed us in using the technology against us. A Cyber Letter of Marque delegating Active Cyber Defense is a small but bold step in changing this paradigm.

CONCLUSIONS

Cyber Letter of Marque is permitted under the United States Constitution and will help deter unrelenting cyber-attacks against the U.S. No treaty or provision of international law prohibits the use of Cyber Letters of Marque by vetted and certified U.S. Corporations.

Jay Healey, senior research scholar at Columbia's School of International and Public Affairs said: "America's cyber power is not at Ft. Meade,…" "NSA and U.S. Cyber Command are simply not positioned, and realistically can't be, to prevent attacks on private sector entities."[28] "By supporting capable businesses seeking to take proactive steps to defend their assets in cyberspace, the new administration can secure a cost-effective policy win with significant potential to improve whole-of-nation cybersecurity."[29]

"Private businesses never anticipated that they would be forced to defend their operations from adversaries as capable as the foreign intelligence services of nation-states. Yet that is what they are forced to do in cyberspace. [T]he American government does not have the resources or bandwidth to be the sole provider of security in this realm. The legal and reputational constraints on the private sector's ability to aggressively and proactively defend itself thus creates a gap in the nation's cyber armor that exposes the integrity of private sector networks and data…"[30]

Leaders of U.S. Corporations are in the best position to quantify how much of their resources to use to defend their own network and assets. In lieu of the taxpayers funding a minimal amount of shared security U.S. Corporations can market their enhanced security and if needed charge for greater protection, thus the market will dictate the amount spent on cyber security and not the federal budget. Similarly, as threats increase, decrease, or change, U.S. Corporations can quickly adjust budgets and deploy the latest technology and personnel much faster than the U.S. Government.

The recent discovery of an epic digital component hardware vulnerability by leading semiconductor companies, and endless software coding errors proves that no hardware or software solution will stop the onslaught of cyber hacking. Only when victims can "hack back" legally and diminish or forestall the gains made by hacking can server farms and networks be more secure.

FRANK COLON is a retired U.S. Navy Judge Advocate and currently a senior civilian cyber operations attorney for the U.S. Army. He has served for over 26 years in various roles with the Department of Defense. He received a bachelor's of business administration in finance from Texas Christian University, a Masters of Arts in National Security and Strategic Studies from the Naval War College, and a Juris Doctorate from Thomas M. Cooley Law School.

ENDNOTES

1 James Lewis, ***Economic Impact of Cybercrime-No Slowing Down***, A Report of the Center for Strategic & International Studies (February 2018) at 4.

2 Rebecca Blumenstein, ***NSA Chief Michael Rogers Talks Cybersecurity,*** Wall Street Journal, November 23, 2016.

3 Simon, ***Raising the Consequences of Hacking American Companies***, at 2.

4 Colonel Gary P. Corn, ***Navigating Gray Zone Challenges In and Through Cyberspace*** (Jan. 16 2018) (2018 Forthcoming) at 6.

5 Id. at 7.

6 CrowdStrike, ***2018 Global Threat Report***, at 73.

7 Simon, ***Raising the Consequences of Hacking American Companies***, at 3.

8 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 47.

9 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 46.

10 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 44.

11 Id.

12 U.S. Const. art. I Section 8 Cl. 11.

13 Eastman, ***Some Famous Privateers*** p. 45 (reproducing a letter of marque granted in 1815 to the ***Grand Turk***).

14 Shock, James R.; Smith, David R., ***The Goodyear Airships***, Bloomington IL, Airship International Press, 2002, p. 43,

15 Richard Struse, Chief Advanced Technology Officer, NCCIC, ***Technical, Policy and Legal Considerations of Cyber Threat Intelligence Sharing***, Department of Homeland Security, (downloaded August 10, 2018: https://www.oasis-open.org/events/sites/oasis-open.org.events/files/1.2%20DHS%20Richard%20Struse.pdf.)

16 Justin Lynch, ***Homeland Security Announces new first response cyber center,*** The Fifth Domain, (August 2, 2018).

17 Mark Rockwell, ***Cyber exercise shows need for closer federal-state coordination,*** FCW, (downloaded February 6, 2019: https://fcw.com/articles/2019/02/06/jack-voltaic-lessons-learned.aspx 2/7/2019).

18 Id.

19 Id.

20 1856 Paris Declaration Respecting Maritime Law (1856), reprinted in ***The Law of Naval Warfare: A Collection of Agreements and Documents with Commentators' 64*** (Natalino Ronzitti ed., 1987) [hereinafter Paris Declaration]

21 Paris Declaration, ***supra*** note 63, at 61-62.

22 Paris Declaration, ***supra*** note 63, at 65.

23 Kessinger, ***Hitting the Cyber Marque***, at 17.

24 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 50.

25 Alexander Melnitzky, ***Defending America Against Chinese Cyber Espionage Through the Use of Active Defenses***, 20 CORDOZO J. INT'L & COMP. L. 537, 540 (2012).

26 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 54.

27 Nye, ***Deterrence and Dissuasion in Cyberspace***, at 52.

28 Sean D. Carberry, ***Why the private sector is key to cybersecurity***, FCW, (March 1, 2017).

29 Frank Cilluffo and Alex Nadeau, ***How the Private Sector Can Remake US Cybersecurity,*** The Daily Signal, (January 31, 2017).

30 id.

11 Robbery or criminal violence at sea.

Need Specialized Technical Support with Easy Contract Terms?

Core Analysis Task (CAT) Program

A Pre-Awarded, Pre-Competed Contract Vehicle.

CSIAC provides Subject Matter Expert (SME) support on an as-needed basis to quickly address technical requirements with minimal contracting effort. CSIAC provides such solutions via the utilization of our Core Analysis Task (CAT) service/capability. CSIAC is a competitively awarded contract with Indefinite Delivery/Indefinite Quantity (ID/IQ) provisions that allow us to rapidly respond to our users' most important needs and requirements. Custom solutions are delivered by executing user-defined and funded CAT projects without the need for further competition.

Through the CAT program, CSIAC is a pre-competed contracting vehicle, enabling the DoD and other agencies to obtain technical support for specific projects/programs that fall within one of the CSIAC technology areas. As with any inquiry, the first four hours are free. If the scope requires a CAT, CSIAC will assist with the development of a Performance of Work Statement (PWS) to be approved by the Contracting Officer's Representative (COR).

Key Advantages of working with CSIAC:

Expansive Technical Domain
The CSIAC's broad technical scope provides numerous pre-qualified resources for potential projects, and is especially valuable for today's information system challenges that frequently cross multiple domains.

Comprehensive STI Repositories
As a consolidation of three predecessor Information Analysis Centers (IACs), CSIAC has a wealth of expertise, data and information to support the successful completion of CATs.

Expansive Subject Matter Expert Network
CSIAC is able to leverage reach-back support from its expansive SME Network, including technical experts from the CSIAC staff, team members, or the greater community, to complete CATs.

Minimal Start-Work Delay
Not only does CSIAC provide DoD and other government agencies with a contract vehicle, but as a pre-competed single award CPFF IDIQ, work can begin in just a matter of weeks.

Apply the Latest Research Findings
CSIAC draws from the most recent studies performed by agencies across the DoD, leveraging the STI holdings of the Defense Technical Information Center (DTIC). The results of all CSIAC CATs and other DoD-funded efforts are collected and stored in DTIC's STI repository to support future efforts by the CSIAC and others.

How To Get Started

If you have a need for CSIAC technical support, the first step is to contact us. All Technical Inquiries are free to the customer for up to four hours of service. If the scope of the support is more extensive and requires a CAT, CSIAC will assist with the development and submission of the task description and related contract documents. CATs may be awarded as either Cost Plus Fixed Fee (CPFF) or Firm Fixed Price (FFP) delivery orders.

Inquiries may be submitted by email to **info@csiac.org**, or by phone at **1-800-214-7921**.

Please visit our website for more information:
https://www.csiac.org/services/core-analysis-task-cat-program/

Who We Are

The Cyber Security Information Systems Information Analysis Center (CSIAC) is the DoD's Center of Excellence in Cyber Security and Information Systems, covering the following technical domains:

- Cybersecurity
- Software Engineering
- Modeling and Simulation
- Knowledge Management/ Information Sharing

CSIAC is chartered to leverage best practices and expertise from government, industry, and academia to solve the most challenging scientific and technical problems. The Center specializes in the collection, analysis, synthesis, and dissemination of Scientific and Technical Information (STI) to produce solutions in support of the defense community.

Our Team

Quanterion Solutions Incorporated is the prime contractor responsible for operating the CSIAC. In addition to Quanterion, customers also have access to the other members of the CSIAC team which include leading technology corporations as well as prestigious academic institutions that perform cutting edge research activities to expand our knowledge base.

CSIAC
Cyber Security & Information Systems Information Analysis Center

266 Genesee Street
Utica, NY 13502

1-800-214-7921
https://www.csiac.org

Transform your Knowledge of the Research Development Test and Evaluation (RDT&E) Budget Process

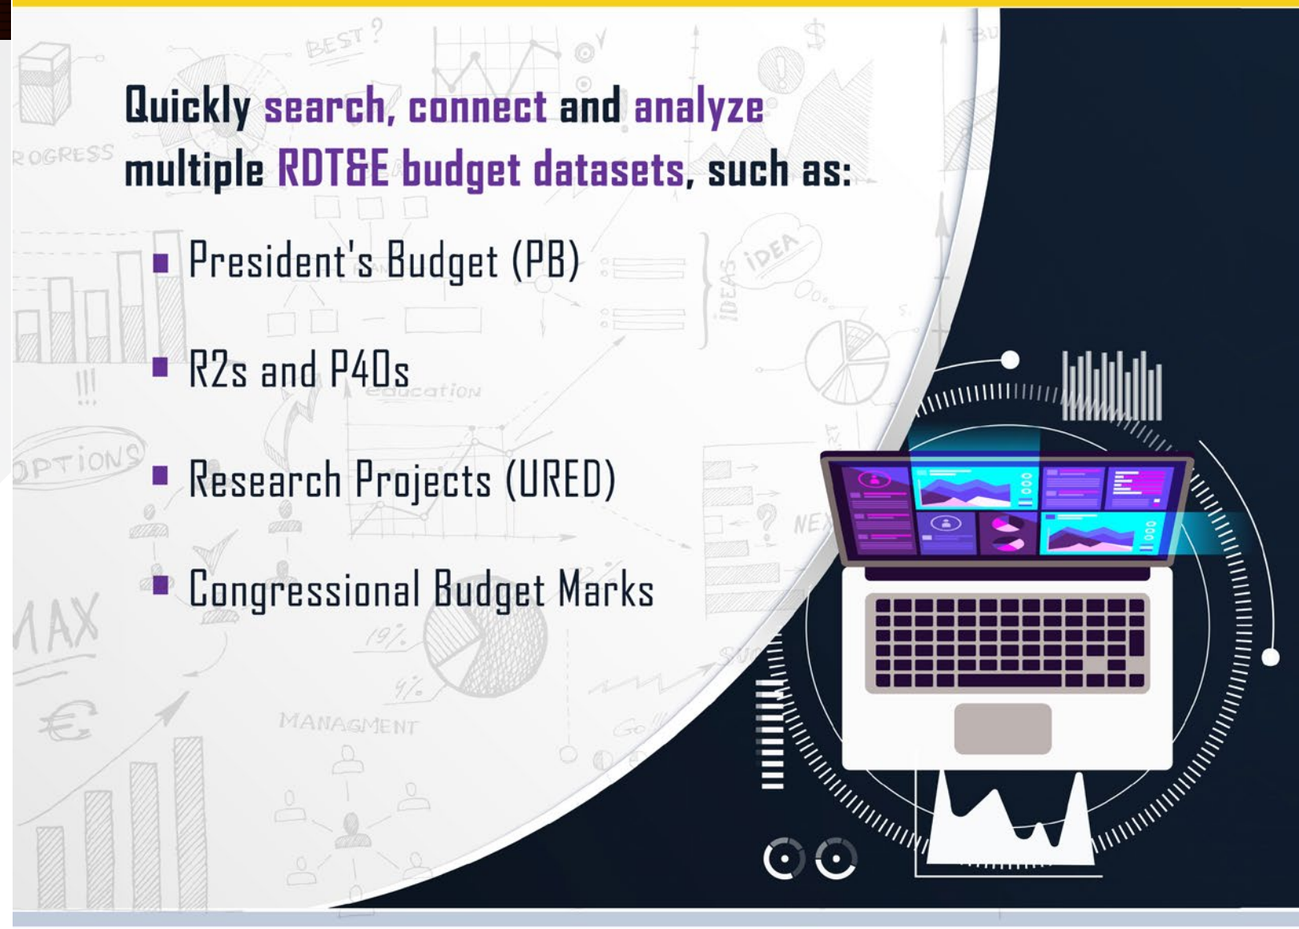

R&E Gateway Powered by DTIC

START TODAY! Visit DTIC' New Research Budget and Project Information (RBPI) tool at https://www.dtic.mil/bt/ui

Defense Technical Information Center (DTIC) | Fort Belvoir, VA | https://discover.dtic.mil

Cyber Security and Information Systems
Information Analysis Center
266 Genesee Street
Utica, NY 13502

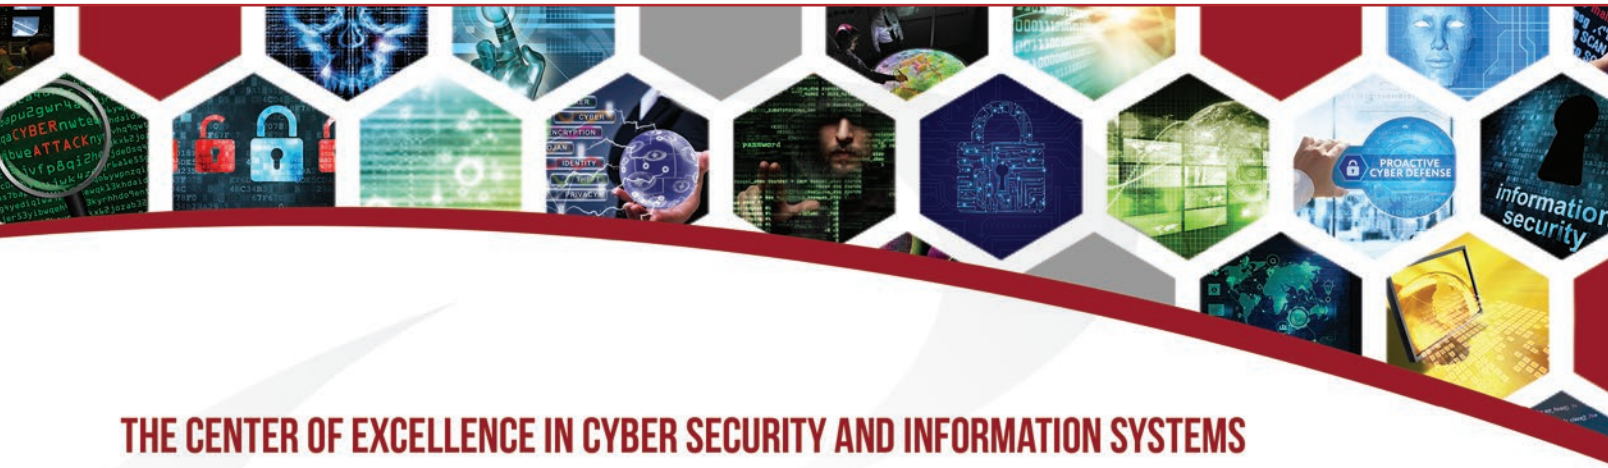

THE CENTER OF EXCELLENCE IN CYBER SECURITY AND INFORMATION SYSTEMS

Leveraging the best practices and expertise from government, industry, and academia in order to solve your scientific and technical problems

https://www.csiac.org/journal/

To unsubscribe from CSIAC Journal Mailings please email us at ***info@csiac.org*** *and request that your address be removed from our distribution mailing database.*