International Journal of Cybersecurity Intelligence & Cybercrime

Volume 5 | Issue 1 Article 4

3-2022

Cybersecurity Risk in U.S. Critical Infrastructure: An Analysis of Publicly Available U.S. Government Alerts and Advisories

cyber threat intelligence, critical infrastructure, information sharing, information extraction, cybersecurity, cyber alerts

IJCIC Center for Cybercrime Investigation & Cybersecurity

Lanz. (2022). *International Journal of Cybersecurity Intelligence and Cybercrime, 5*(1), 43-70.

Cybersecurity Risk in U.S. Critical Infrastructure: An Analysis of Publicly Available U.S. Government Alerts and Advisories

Zachary A. Lanz*, University at Albany, State University of New York, U.S.A.

Keywords: cyber threat intelligence, critical infrastructure, information sharing, information extraction, cybersecurity, cyber alerts

Abstract:
As threat actor operations become increasingly sophisticated and emphasize the targeting of critical infrastructure and services, the need for cybersecurity information sharing will continue to grow. Escalating demand for cyber threat intelligence and information sharing across the cybersecurity community has resulted in the need to better understand the information produced by reputable sources such as U.S. CISA Alerts and ICS-CERT advisories. The text analysis program, Profiler Plus, is used to extract information from 1,574 U.S. government alerts and advisories to develop visualizations and generate enhanced insights into different cyber threat actor types, the tactics which can be used for cyber operations, and sectors of critical infrastructure at risk of an attack. The findings of this study enhance cyber threat intelligence activities by enabling an understanding of the trends in public information sharing as well as identifying gaps in open-source reporting on cyber-threat information.

Introduction

Cyber threat actors continuously demonstrate an increasing capability for sophisticated attacks, putting critical infrastructure at a high risk of falling victim to malicious operations, whether financially or strategically motivated. In response to this growing risk, cybersecurity professionals across the industry have emphasized the need for better information sharing, and for actionable intelligence to mitigate the risk of a potential attack. As the cyber risk landscape is constantly evolving, equipping the cybersecurity industry with knowledge of the trends and patterns at the root of open-source reporting can provide analysts with a new opportunity to drive threat intelligence activities, predict threat behavior, and work to mitigate risks through information sharing. This is separate and apart from the ability of IT security specialists to leverage information sharing and intelligence sharing to simply patch or remove vulnerable software.

Public reporting on cyber threats has been growing throughout the past decade. However, there has yet to be a comprehensive analysis of the information disclosed in such reports and the trends and patterns hidden within them. An analysis of publicly available reporting on cyber threat actors, the techniques they employ for operations, and the sectors of critical infrastructure at risk of an attack, can provide the cybersecurity community with a stronger understanding of the publicly available information at their disposal (as it relates to critical infrastructure cybersecurity). The findings derived from such an analysis can be used to enhance both cyber practitioners' and academics' understanding of the content and trends present in government cyber reporting while also identifying gaps in the public information sharing landscape. The goal of this study is to analyze publicly available reporting on cyber threats and develop an understanding of the changes in reporting, throughout the past decade. This thesis is divided into five distinct sections: literature

*Corresponding author
Zachary A. Lanz, University at Albany, 1400 Washington Ave, Albany, NY 12222, U.S.A.
Email: zacharylanz@live.com

on the topic; 2. the approach taken to conduct the study; 3. an analysis of the results; 4. implications for future research; and 5. conclusions drawn from the results.

Literature Review

To effectively analyze the information provided in public cyber reporting, it is first necessary to develop a comprehensive outlook on existing information surrounding cyber threat actors, cyber threats to critical infrastructure, and the techniques attackers have employed to conduct malicious operations. By conducting a review of these focus areas, a foundation of information can be built and leveraged to inform which key concepts in the cybersecurity domain require the most attention during an analysis of public reporting. In the last two decades, a variety of cyber threat actors have emerged seeking to exploit a range of technologies, to threaten the United States national security interests, and to advance their own strategic and economic objectives (Campbell, 2018; Coats, 2019; Daş & Gündüz, 2019). The growth and increased reliance on new connected technologies, particularly for critical infrastructure, has generated a new risk landscape. The integration of internet-connected devices into infrastructure environments has resulted in several attacks against critical services, such as the Ukrainian Electric Grid Compromise of 2015 and an attack targeting a small dam in New York in 2013 (see Malgaras et al., 2018). As demonstrated by a recent surge in ransomware attacks against hospitals, the ability of threat actors to cause immense damage to U.S. infrastructure poses an increasing risk to public health, safety, and prosperity as new technologies are integrated into infrastructure, networks, and services (Cimpanu, 2020; Coats, 2019; Landi, 2020).

Nations, criminal groups, terrorists, hackers, hacktivists, and insiders, all pose significant threats to critical infrastructure, which often rely on Supervisory Control and Data Acquisition (SCADA) and Industrial Control Systems (ICSs) (CISA, n.d.; Fleury et al., 2008; GAO, 2019; Malgaras et al., 2018). These actors are targeting different critical infrastructure sectors, with threats to one ICS entity driving threats to other industry verticals (see Dragos, 2020; Crowdstrike, 2020). Attacks targeting critical infrastructure can lead to physical damage, injury, environmental effects, or even loss of life (Malgaras et al., 2018). As cyber threat actors become increasingly capable of targeting and compromising critical infrastructure environments, the likelihood and severity of a cyber-attack will increase (Ahmad et al., 2019; Jang-Jaccard & Nepal, 2014).

Each of the commonly identified threat actor types (nations, criminal groups, terrorists, etc.) has been associated with various motivations for conducting their operations. For instance, nation-state sponsored groups are often connected to operations involving information gathering, espionage, and the ability to cause service disruptions for strategic purposes, whereas cyber-criminals are commonly driven by financial gain, and terrorists may seek to send coercive political messages (GAO, 2019; Crowdstrike, 2020; CISA, n.d.). Hacktivists, hackers, and insiders break into networks for a variety of reasons, ranging from revenge to monetary gain (CISA, n.d; Fleury et al., 2008; GAO, 2019). Internet of things (IoT) technologies and open-source information on the web have made it much easier for actors to carry out malicious cyber-enabled activities. Insiders can inflict damage to an organization simply by using their company-issued laptop with ill-intent, and terrorists and spies can use open-source intelligence (OSINT) for most of their information-collection needs (Geers, 2009; GAO, 2019).

The capabilities of cyber threat actors were demonstrated in the Ukraine Electric Grid Compromise of 2015, when the Sandworm Team (widely associated with Russian military intelligence), caused power outages by

hacking into a Ukrainian energy company's SCADA system to advance Russian geopolitical interests (E-ISAC, 2016; Nussbaum, 2019). The Sandworm Team is one of many highly capable Advanced Persistent Threat (APT) actors, a term often used in reference to highly motivated, sophisticated, and organized threat groups (Ahmad et al., 2019; Lemay et al., 2018). Unfortunately, various challenges arise when tracking APTs. One of the most prominent complexities includes the lack of a widely adopted ontology in the cyber domain; an example of which can be found in how organizations identify and track the same APTs under different names (Lemay et al., 2018; Mavroeidis & Bromander, 2017). The lack of formal standards surrounding cyber knowledge management is a significant limitation on effective (and necessary) information sharing (Dandurand & Serrano, 2013; Mavroeidis & Bromander, 2017). Russia is not the only country conducting largescale cyber operations; China, Iran, and North Korea are also recognized to be playing prominent roles in threatening United States national security with their advancing cyber capabilities (Abdyraeva, 2020; Coats, 2019).

Evolving Tactics in a Shifting Threat Landscape

Threat actors employ what is known as tactics, techniques, and procedures (TTPs) to conduct their operations. These TTPs characterize an adversary's behavior in terms of what they are doing and how they are doing it (Mavroeidis & Bromander, 2017). TTPs, at their core, are the methods threat actors use to orchestrate and manage their attacks, and thus an understanding of these activities is an essential component of effective cyber threat intelligence (CTI) (Dunham, 2017; Mavroeidis & Bromander, 2017). An analysis of the various TTPs that adversaries use can enhance CTI activities by supporting threat modeling, triaging, contextualization, and attack vector identification (Dunham, 2017; Lemay et al., 2018). TTPs help attribute attacks to actors by mapping the actions taken (in an attack), generating attack signatures, and identifying similarities with previous events (Dunham, 2017; Tounsi & Rais, 2018).

There are many distinct TTPs, from ransomware to phishing scams, or remote code executions, to the use of botnets (MITRE, n.d.; Mavroeidis & Bromander, 2017; SecureWorks, 2017; Crowdstrike, 2020). However, there are currently no widely adopted classification schemes for the various TTPs. Several scholars have developed taxonomies to classify cyber-attacks using a range of different models (Bahrami et al., 2019). Some of these taxonomies have followed the Cyber Kill Chain (CKC) model, mapped the varying stages of an attack through the MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) framework, or have even generalized categories of TTPs based on the most common successful attack types (Bahrami et al., 2019; MITRE, n.d; Mavroeidis & Bromander, 2017). Research on the development of CTI-driven ontologies reported that a vaguely defined terminology can lead to significant confusion among cyber experts, expressing the need for a standardized knowledge structure in the cyber domain (Mavroeidis & Bromander, 2017).

In addition to the range of tactics and techniques at a threat actor's disposal, cyber-attacks have been occurring at a higher frequency, with the TTPs employed becoming increasingly sophisticated (Abdyraeva, 2020; Dion et al., 2009). Threat actors employ complex TTPs including using advanced anti-forensics and evasion methods in their attacks to prevent detection (Conti et al., 2018). The advanced capabilities and sophistica sophisticated tactics demonstrated by threat actors go beyond exploiting primary targets in their ability to compromise critical infrastructure through a range of different attack vectors (Checkpoint, 2020; Mavroeidis & Bromander, 2017). Capable actors use a broad range of attack techniques, cyber and/or

physical, and constantly evolve their tactics to overcome organizational defense mechanisms (Bahrami et al., 2019). To meet this challenge against critical infrastructure head-on, the use of CTI is necessary to develop tactical, operational, and strategic knowledge of threat actor activities and employed TTPs (Oosthoek & Doerr, 2020; Work, 2020).

Critical Infrastructure at Risk

In 1997, the first major policy document regarding critical infrastructure protection (CIP), the President's Commission on Critical Infrastructure Protection (PCCIP), was released. This report reflected the United States' growing fear of attacks against domestic assets and critical services (Dion et al., 2009). Since then, numerous CIP-related offices have been created across federal, state, and local levels; and the recent increase of IoT applications in ICS environments has generated a new emphasis on the protection of infrastructure through the cyber domain (Dion et al., 2009; Daş & Gündüz, 2019). In response to the growing risk of cyber-attacks against critical sectors, entities including the ICS Cyber Emergency Response Team (ICS-CERT) and the Cybersecurity and Infrastructure Security Agency (CISA), both of which fall under the Department of Homeland Security (DHS), were created (see CISA, n.d.; Greenberg, 2019; Weiss, 2010, pp. 74-97). CISA was formed in 2018, drawing on the previous National Protection and Programs Directorate (NPPD) within DHS, to protect the nation's critical infrastructure from physical and cyber threats with the potential to cause significant damage(s) (Brumfield, 2019; Daş & Gündüz, 2019). Governments and private organizations have created industry-specific CTI "sharing groups," known as information sharing and analysis centers (ISACs) to mitigate risk by communicating intelligence on threat actor campaigns, motivations, indicators of compromise (IoCs), and vulnerabilities, driving a comprehensive outlook into relevant threats (Koepke, 201&; Ring, 2014; Wagner et al., 2019).

In recent years, numerous cyber-attacks targeting ICS and SCADA systems have been identified, resulting in cybersecurity becoming one of the main concerns for SCADA and ICS operators (Daş & Gündüz, 2019; Maglaras et al., 2018). In conjunction with the targeted cyber-attack against Ukraine's Energy Sector in 2015, the attackers also targeted Ukraine's Government, Communications, and Transportation Sectors (E-ISAC, 2016; Layne, 2017). A few years before the Ukraine incident, in 2010, a targeted cyber-attack occurred against Iran's nuclear sector. That attack targeted the Natanz Uranium enrichment facility and caused physical destruction to the centrifuges onsite by altering their rotational speed (Layne, 2017; Karnouskos, 2011). More recently, several Israeli Water Authority facilities were targeted in what is widely believed to be an attack by Iranian threat actors (Olenick, 2020). Threat actors have become increasingly capable against infrastructure around the world, driving the need for proactive ICS cybersecurity and the use of CTI to mitigate the increasing number of cyber-attacks against critical infrastructure (Gritzalis et al., 2019, p. 221-224; Ring, 2014; Wagner et al., 2019; Weiss, 2010).

Each of the 16 critical infrastructure sectors identified by the Presidential Policy Directive 21 (PPD-21) and defined by CISA is considered vital to the U.S.; their incapacitation or destruction would have severe effects on the nation (CISA, 2020; Smith, 2015). The 16 sectors of critical infrastructure defined by CISA: Chemical; Commercial; Communications; Critical Manufacturing; Dams; Defense Industrial Base; Emergy; Financial Services; Food and Agriculture; Government Facilities; Healthcare and Public Health; Information Technology; Nuclear Reactors, Materials, and Waste; Transportation Systems; and Water and Wastewater Systems. However, the energy sector plays a unique role in the national framework, as it provides an

"enabling function" across all critical infrastructure sectors (CISA, 2014). The energy sector poses a significant risk to national security due to external reliance on its services as well as an increasing number of reports about adversaries targeting the U.S. electric power system (Campbell, 2018). In May of 2020, U.S. President Donald J. Trump signed an executive order to secure the country's bulk-power system against exploitation and attacks by foreign adversaries (Department of Energy, 2020). This order reflects the continued concern of threats posed by capable cyber-adversaries against the United States. As the number of attacks against ICS environments increases globally, so does the risk of a disruptive cyber event against U.S. critical infrastructure, furthering the need to mitigate risks through information sharing practices and CTI initiatives (Dragos, 2020; Tounsi & Rais, 2018; Wagner et al., 2019). To develop an understanding of the threat actors targeting critical infrastructure and the TTPs they employ to conduct operations, a study surrounding the contents of publicly available government reporting can take place.

Data and Methods

An analysis of public reporting on TTPs, critical infrastructure sectors, and adversary types can drive threat intelligence activities by supporting threat trend modeling, attack contextualization, and even providing insight into changes in CTI reporting patterns (Dunham, 2017; Lemay et al., 2018). This type of information is currently disseminated in alerts and advisories issued by the DHS CISA and ICS-CERT but is not aggregated into a data set suitable for CTI activities. The development of such a dataset can provide a direct understanding of the trends and patterns across publicly available U.S. Government cyber reporting. In addition, the results can deliver a new outlook on the U.S. Government's perception of cyber threats and changes in the type of information they are willing to publicly disclose. Other federal agencies publish cyber-related alerts. However, DHS CISA and ICS-CERT provide the most consistent and publicly available reporting that discusses TTPs, critical infrastructure, and adversaries.

Although the alerts and advisories provide different information, they are both designed to deliver timely and reliable intelligence on current security issues, vulnerabilities, and exploitation activities (CISA, n.d). The CISA alerts, which are produced through the National Cyber Awareness System (NCAS), disseminate actionable CTI surrounding the activities of threat groups, malware strains, vulnerabilities, and recommendations for securing various types of technology and infrastructure. Similarly, the ICS-CERT advisories, which are generated specifically for ICS operators, provide information on vulnerabilities that impact products in ICS environments. These advisories also detail how attackers can exploit weaknesses and determine which critical infrastructure sectors are at risk or impacted by the disclosed vulnerability or threat.

CISA and ICS-CERT publish reporting on more than just NCAS alerts and vulnerability advisories. However, these two publication types provide comprehensive details on threat actor activities, TTPs, and impacted sectors of infrastructure. The information provided in these alerts is not reflective of all cyber-attacks against critical infrastructure but is reflective of public government reporting on attacker activities, vulnerabilities that enable adversary operations, and relevant TTPs, providing information that can be used to understand possible trends. To this end, the CISA Alerts are used to draw insights into the types of threat actors commonly reported, whereas the ICS-CERT Advisories are used to understand the sectors of infrastructure most impacted by reported vulnerabilities and possible exploitation activities. However, both publication types are used to gain an understanding of the TTPs relevant to attacks and potential exploits.

Approximately 1,371 ICS-CERT advisories dating back to early 2010, and 203 CISA Alerts dating back to 2008 are analyzed to achieve this goal. The data set for this study includes all alerts and advisories from November 2008 until June 2020 as that is when the most up-to-date documents were obtained. CISA also provided alerts that disclosed information surrounding vulnerability updates from 2004 through 2007. However, these alerts were identified as archived, no longer updated, and potentially contained outdated information; thus, they were not included in the study. The alerts and advisories were manually downloaded as text files from the CISA website and were individually reviewed to ensure the full content of each document was included in the download.

A Rule-Based Information Extraction Approach

A rule-based information extraction (IE), named the Cyber Alert Information Extraction Scheme (CAIES), was developed to analyze, collect, and output data from the alerts using Profiler Plus. CAIES is an original scheme developed for the sole purpose of conducting this study and specifically designed to target the cybersecurity-related terminology present in the ICS-CERT advisories and CISA alerts. Profiler Plus is a general-purpose natural language analysis application that allows for multi-pass, rule-based analysis of text, and relies on researcher input and specification rather than machine learning (ML) (Levine & Young, 2014; Neuendorf, 2018). Employing a rule-based IE approach provides a range of benefits. Rule-based IE is easy to comprehend, easy to maintain, and easy to infuse with domain knowledge (Chiticariu et al., 2013). Additionally, rule-based systems can be considered compositional as they are adjustable and make it easy to trace and rectify the causes of any errors that may arise during development (G.C. et al., 2015; Waltl et al., 2018). However, rule-based IE systems are heuristic in nature and are known to be very time-consuming as they require tedious manual labor (G.C. et al., 2015). Although these challenges were present throughout the development cycle, a thorough understanding of the context and content of the alerts and advisories were needed to incorporate domain-related knowledge into the study, turning the heuristic and time-consuming nature of rule-based IE systems into an advantage for more informed and accurate results.

An alternative to a rule-based system is the statistical ML approach. ML-based approaches are trainable, adaptable, and have become the most widely adopted IE approach in modern academic research (Chiticariu et al., 2013; Waltl et al., 2018). However, ML techniques require deep theoretical knowledge to use or maintain and require retraining for domain-specific adaptation (Chiticariu et al., 2013; Waltl et al., 2018). Large amounts of labeled datasets are needed when employing an ML-based approach, requiring a significant amount of manual labor. Additionally, while rule-based systems only require the adjustment of a few rules to reconfigure targeted information or make adjustments, ML-based approaches require complete retraining of the data (Chiticariu et al., 2013). The importance of incorporating domain knowledge for effective information extraction and the ease of adaptability in rule development made rule-based IE the optimal approach for conducting this study.

Rule and Dictionary Development

The rules developed for this IE rely on newly developed dictionaries that identify key terms and classify the specific alert and advisory number along with its original release date. To effectively detect mentions of 1) threat actors, 2) TTPs, and 3) involved sectors of critical infrastructure, three dictionaries, and three outputs were developed for each category. These three dictionaries are essential to properly identify the

correct terms that need to be detected and marked for the final outputs. However, as there is no formal ontology or widely adopted standard for cybersecurity definitions, dictionary development proved to be challenging (Conti et al., 2018).

The ThreatTable dictionary. The "ThreatTable" dictionary is designed to detect mentions of different threat actor (Adversary) types. The foundation for this dictionary was CISA's classifications and descriptions of cyber threat sources, which identifies national governments, terrorists, industrial spies and organized crime groups, hacktivists, and hackers as primary threat types (CISA, n.d.). Although this provided a strong basis for the ThreatTable Dictionary, the CISA classifications are not consistently used in the alerts and advisories, requiring some additional classification of threat actor types based on manual searches across the publications. As the aim of this study is to identify the various types of threat actors being mentioned in public reporting, a deep dive into specifically attributed adversaries will not occur. Even though there is mention of specific state-sponsored or cyber-criminal groups, discussions surrounding which groups or nations are mentioned fall outside the scope of this analysis.

The CISectors dictionary. The CISectors Dictionary identifies words and phrases that correspond to the 16 sectors of critical infrastructure (classified as CI-Sectors) defined by CISA (CISA, 2020). This dictionary required little to no revision as both the ICS-CERT advisories and CISA alerts consistently identified the same 16 sectors using largely similar terminology. The only additions made to the dictionary that are not the same as CISA's defined sectors of critical infrastructure were variant terms referring to some of the same sectors. For example, some alerts refer to the "energy sector" as the "electric grid." In addition, several early publications included "building automation" as an infrastructure sector. However, this is no longer acknowledged as a sector and is not included in the output due to its extremely low hit count during preliminary analysis.

The TTP dictionary. The term TTP is broad and can include anything from the type of malware a threat actor uses to their strategic objectives (i.e., financial information/data theft) or attack methodology (i.e., privilege escalation) (Dunham, 2017; Conti et al., 2018). TTPs are not mutually exclusive, and several can be used in a single attack. Without a widely adopted ontology for common TTPs, the TTP Dictionary was created based on open-source information provided by a variety of cybersecurity-focused companies and public organizations (see, Dunham, 2017; Secureworks, 2017; MITRE, n.d.). The preliminary analysis of the publications provided insight into some of the more common TTP types discussed within the alerts, indicating which TTPs would be the most prevalent and/or useful to include in the output.

Rule-Based IE Performance Evaluation

Academic papers evaluate the performance of an IE system in terms of precision and recall over standard labeled data sets in comparison to peer systems (see Akhtar et al., 2020; Antoun et al., 2020; Lybarger, 2020). To evaluate the accuracy, precision, and recall of the coding scheme developed for this study, 48 of the alerts were fully annotated (See Appendix A for Annotation Guide). Additionally, a set of rules were created to identify and count the occurrences of True Positives, False Positives, True Negatives, and False Negatives across the annotated documents (although True Negatives are removed as they skew the accuracy of the results) (see Koehrsen, 2018). The performance evaluation includes the following:

Table 1. Rule-Based IE Performance Evaluation

True Positives	738	False Positives	122
False Negatives	43	Accuracy	82%
True Negatives	121736	Precision	86%
		Recall	94%

One of the greatest takeaways from these results is that recall - the ability of the model to identify all relevant data points - received a high success rate of 94% (Koehrsen, 2018). While recall represents the coding scheme's ability to identify all relevant instances, precision represents the proportion of the data that is relevant (Koehrsen, 2018). Both measures are valuable to the goal of identifying accurate and relevant occurrences of threat actors, TTPs, and CI-Sectors. These results demonstrate the coding scheme's ability to achieve that goal. The coding scheme is also able to accept noise that leads to lower accuracy and precision due to the high recall - the most critical measure for depicting the coding scheme's ability to identify the presence of the dictionary terms within the alerts (Weischedel & Boschee, 2018). The accuracy of the coding scheme is not as high as its recall (due to the presence of false positives and false negatives). However, a manual review of the category outputs shows that, at the observation level, the coding scheme accurately reports the presence of each category within an alert 100% of the time, marking the true success of the IE. Although a baseline for performance cannot be set based on peer systems, as no peer systems exist, success can be measured by the meaningfulness of the data extracted, particularly at the observation level.

Results and Discussion

To effectively analyze the results, the output from Profiler Plus was imported into the data visualization software Tableau. Tableau provided a direct means to create graphical representations of the data extracted from the alerts and advisories; an accessible way to drive insight into the trends and patterns within the datasets (Tableau, n.d.). Tableau offers a variety of data visualizations, anything from text charts and bar charts, to line graphs, area charts, and more (Tableau, n.d.). These visualizations were used to achieve the study's goal of developing an understanding of the trends and patterns across public U.S. Government cyber reporting.

Figure 1 demonstrates how public reporting from CISA's NCAS and ICS-CERT has generally increased from 2008 to early June 2020. Alerts have been published through NCAS since as early as 2004, but the alerts from 2004 to 2007 were archived by CISA and were not included in this study. It is evident that there has been a significant increase in the public disclosure of cyber threat-related information throughout the past decade. Based on the 145 alerts released within the first six months of the year, reporting at the end of 2020 will be similar to that of 2018 or 2019, or higher. The increasing release of alerts each year, as shown in Figure 1, can be expected to continue in the coming years based on the pattern of growth since 2008. Figure 1 potentially indicates a growth in the willingness of the U.S. Government to provide public information surrounding cyber-related threats and activities. It is also likely that the increase in reporting throughout the past decade is in response to a growing demand for stronger information-sharing throughout the cybersecurity community (Conti et al., 2018; Dandurand & Serrano, 2013).

Figure 1. Total Number of CISA Alerts & ICS-CERT Advisories Per Year

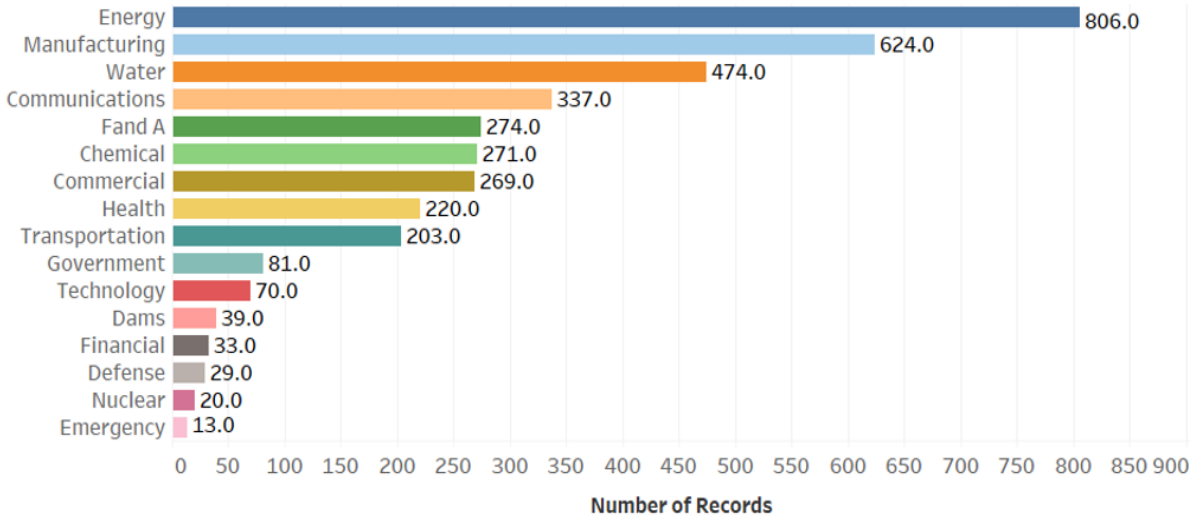

Figure 2. Top Reported CI-Sectors Across CISA Alerts and ICS-CERT

Of the reported sectors of critical infrastructure, Energy, Manufacturing, Water (and Wastewater), and Communications were, as anticipated, identified as the most mentioned sectors across all the alerts and advisories. On the other hand, Financial Services, Defense Industrial Base, Nuclear, and Emergency Services sectors were the least mentioned sectors across the alerts. Although some sectors are more prevalent within the alerts than others, Figure 2 indicates that all sectors have been identified as "at-risk"

of a form of cyber-attack, major vulnerability, or threat. The low number of reports discussing the Defense, Nuclear, and Emergency Services sectors may be due to the sensitive nature of those sectors.

Figure 3 provides a timeline of reporting segmented by sector (see Appendix B for sector-specific timelines). Although there had been a low level of reporting throughout 2008 and 2009, 2010 marked the start of a rapid increase in the number of alerts released that specifically mentioned different sectors of critical infrastructure. Most notably, there is a large spike in energy sector-related reporting in 2010, something that coincides with the discovery of Stuxnet. Stuxnet changed how the world views the implications of cyber-attacks against critical infrastructure with its ability to tamper with Iran's nuclear program, resulting in a newfound focus on ICS security (Greenberg, 2019; Zetter, 2016). This first-of-its-kind cyber-weapon marked the beginning of a new type of warfare in the cyber realm. Stuxnet set a precedent for attacks against infrastructure and demonstrated how challenges of attacker attribution could minimize the potential for the same geopolitical fallout that would result from a kinetic attack (Zetter, 2016).

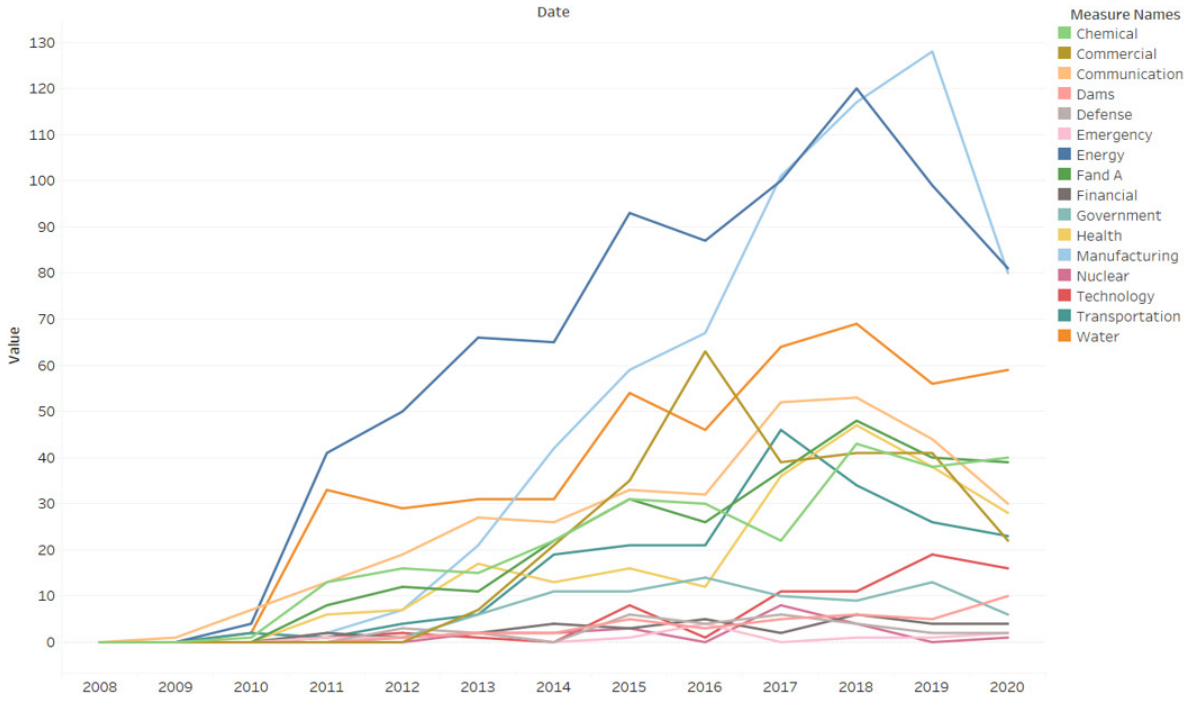

Figure 3. CI-Sector Specific Reporting Timeline Across CISA Alerts and ICS-CERT Advisories

The Energy sector was mentioned in approximately 51% of the total alerts analyzed. This sector, in particular, is a major concern when considering the possible impacts of a cyber-physical attack, especially since witnessing the capabilities of cyber-weapons as demonstrated by Stuxnet in 2010 (Zetter, 2016). In 2011, Pentagon officials declared that cyber-attacks against the U.S electric grid would be considered an act of war, and would warrant an appropriate response, including military options (Zetter, 2016, p. 398). The

emphasis on energy reporting since 2010 directly reflects the growing cyber risk across the energy sector, and concerns over energy cybersecurity have further increased following the compromise of Ukraine's electric grid in 2015 (Greenberg, 2019). Stuxnet set a precedent for cyber-attacks against critical infrastructure, and similarly, the Sandworm Team proved that an entire region of electricity could be taken offline through the exploitation of undisclosed zero-day vulnerabilities (Greenberg, 2019, p. 100).

Adversary Analysis

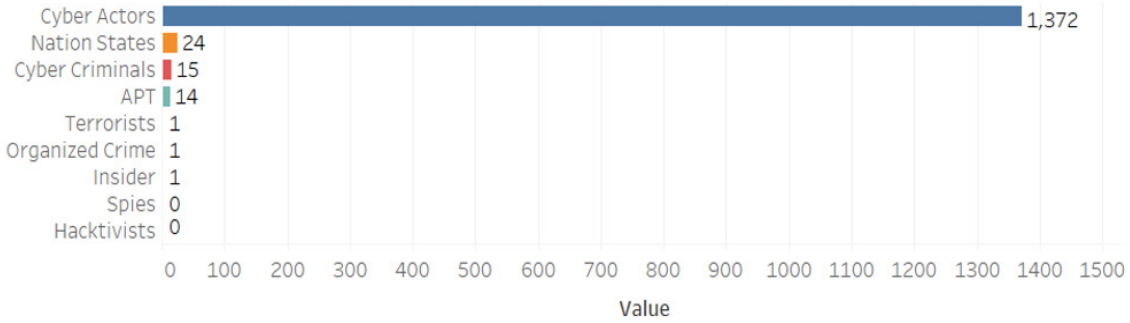

Figure 4. Top Adversary Types Reported Across CISA Alerts and ICS-CERT Advisories

Figures 4, 5, and 6 provide a unique insight into the changes in adversary attribution throughout the past decade. As depicted in Figure 4, the most common form of reporting on adversaries was that which used generalized terminology (i.e., Cyber-actors, threat actors, attackers, hackers, etc.) and was not specific to the identity of the adversaries involved. These generalized attributions accounted for approximately 87% of the total alerts analyzed. When considering the content of the ICS-CERT advisories, which comprised 87% of the total alerts, a lack of adversary attribution is to be expected as the advisories are typically used to communicate vulnerabilities and possible exploitation techniques rather than threat actors' behavior and activities. Historically, threat actor attribution has posed a significant challenge in the cybersecurity industry as the hackers behind any operation, especially sophisticated ones, are often impossible to specifically identify (Eichensehr, 2020; Greenberg, 2019, p.13).

Although low in number compared to the entire sample of alerts analyzed, many mention nation-states actors, APT groups, and cyber-criminals. When focusing on these three adversary types (as seen in Figure 5), we see that the first signs of public attribution to a specific adversary type occurred in 2011 and that following 2016, there was a rise in attribution to specific adversaries compared to previous years. Mention of terrorists and insiders occurred for the first time in 2020.

The shift towards more specific adversary attribution within public reporting may reflect both a direct change in attribution capability as well as a shift in U.S. policy. Figures 5 and 6 show a slight increase in nation-state attribution following 2016, along with an increase in APT reporting. Coincidentally, 2016 marked the first time the FBI and DHS directly attributed malicious cyber activity to specific countries within a joint analysis report (JAR) (NCCIC & FBI, 2016).). This first-of-its-kind JAR was a result of analytic efforts between the FBI and DHS which disclosed intelligence into an ongoing Russian cyber-enabled operation aimed at interfering with the 2016 U.S. presidential elections (NCCIC & FBI, 2016).

Although the 2016 JAR was the first to publicly attribute a cyber-campaign to a nation-state, it was by no means the last time such attribution would occur. The U.S. Department of Justice (DOJ) has indicted several individuals tied to cyber-campaigns since 2016 (see Eichensehr, 2019). It is also important to note that the lack of specific adversary attribution is likely a direct result of the sensitivity surrounding threat intelligence. The data being analyzed is that which is publicly reported by U.S. authorities. This does not entirely reflect the knowledge and capabilities of such authorities but rather the information that they are willing to publicly disclose.

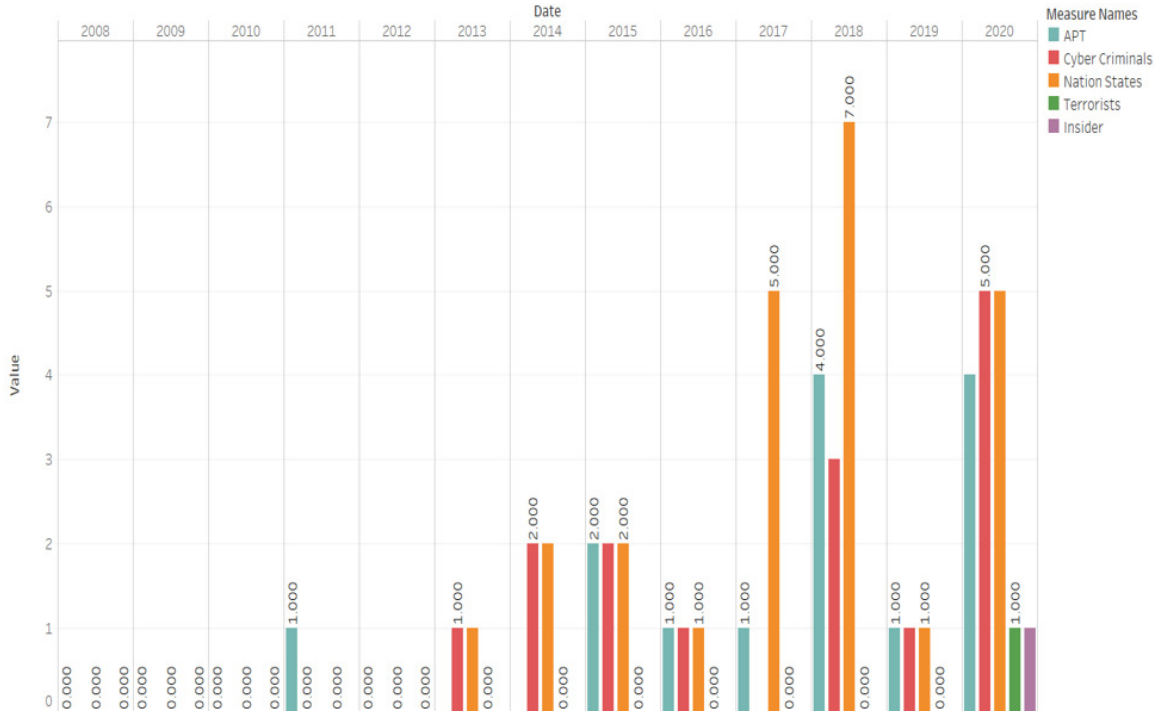

Figure 5. Top Attributed Adversary Types Within CISA Alerts and ICS-CERT Advisories

Unfortunately, the efficacy of deterring adversaries through public attributions (i.e., indictments, press releases, technical alerts) may not be as successful as one would hope. Although public attribution may generate some level of deterrence, by no means does it entirely prevent the occurrence of cyber-attacks (see Eichensehr, 2019). Also, governments may be disincentivized to publicly attribute attacks to specific nation-state groups as they may face pressure to undertake punitive measures (Eichensehr, 2019). This is vastly different from non-governmental entities who might attribute specific governments to cyber-activity as they are not obligated or expected to take responsive action (Eischensehr, 2019). Governmental and non-governmental entities who attribute actors to other countries can make themselves an unwanted target for retaliation (Edwards et al., 2017; Eichensehr, 2019, p. 24). On top of possible sensitivity issues (i.e., classified intelligence) the strategic and technical challenges behind attribution to specific adversaries, especially those state-sponsored, likely play a role in the overwhelming amount of generalized attribution seen in the results (Figure 6).

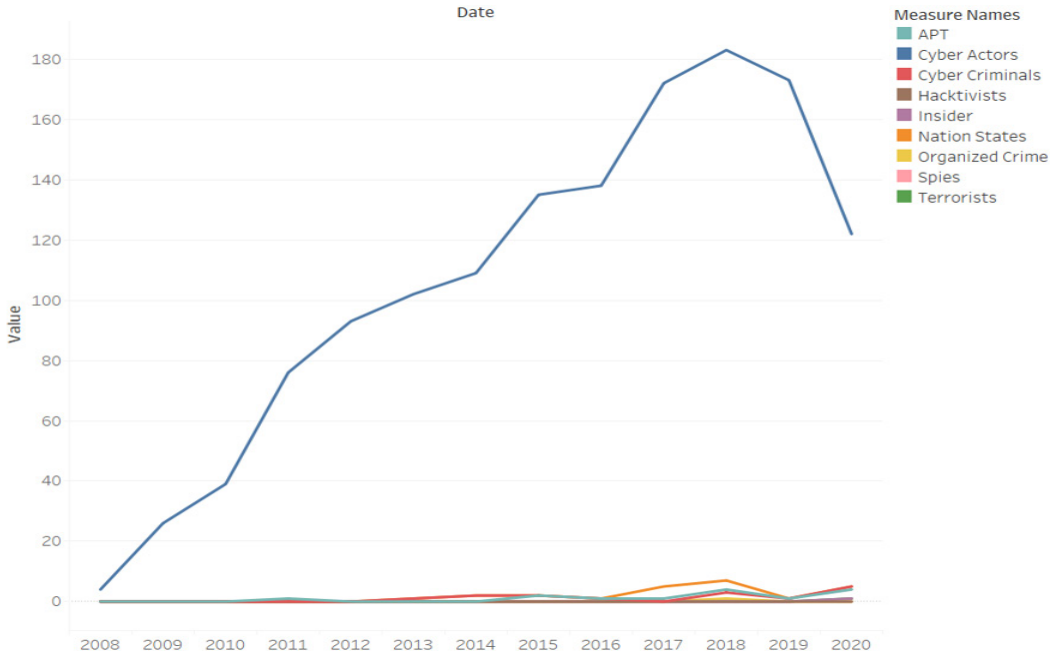

Figure 6. Adversary Specific Reporting Timeline

TTP Analysis

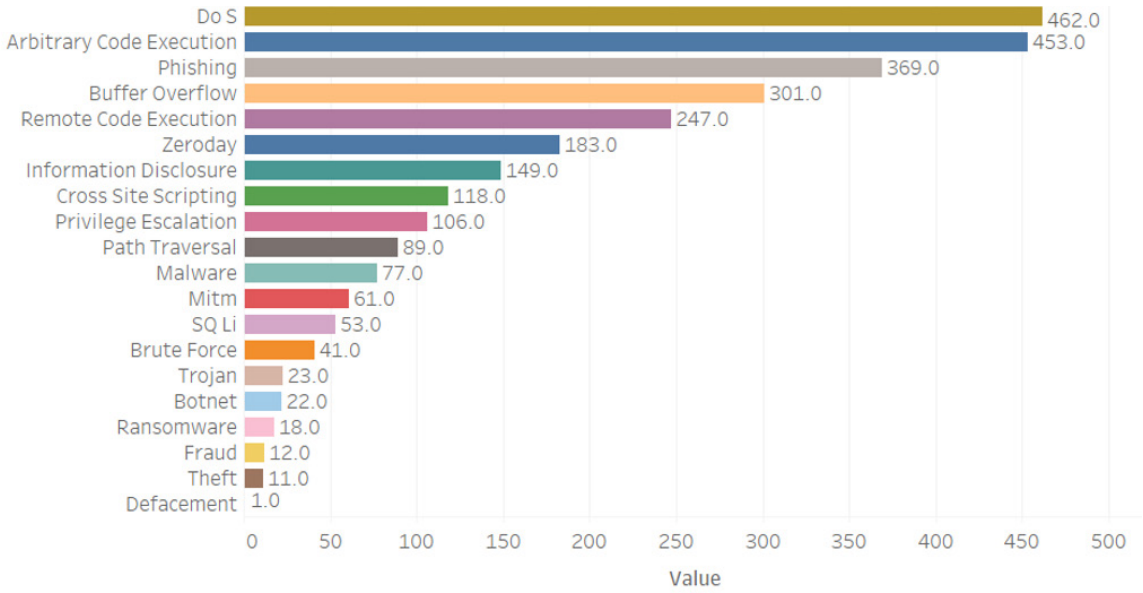

Figure 7. Top TTPs Reported Across CISA Alerts and ICS-CERT Advisories

Figure 7 shows the results of the TTP output, visualizing mentions of the TTPs defined in the TTPs dictionary (see Appendix C for full TTP reporting timelines). Although TTPs are not mutually exclusive and can often drive one another throughout an attack cycle, the top results generated through this analysis appear to be somewhat distinct from one another. The only exception to this is buffer overflow-related attacks, which often work to enable remote or arbitrary code execution and denial of service (DoS) attacks. DoS attacks have been a prevalent threat for years and have been used by a variety of threat actors to disrupt the availability and operation of a function or resource, impacting the ability of an organization to provide a specific service (Jang-Jaccard & Nepal, 2014). High results for DoS, code execution attacks, and phishing are unsurprising due to their prevalence in existing literature (see Bendovschi, 2015; Bitdefender, 2020; Salahdine & Kaabouch, 2019). However, the lower numbers in ransomware reporting and even malware is entirely unexpected (Figures 7 & 8).

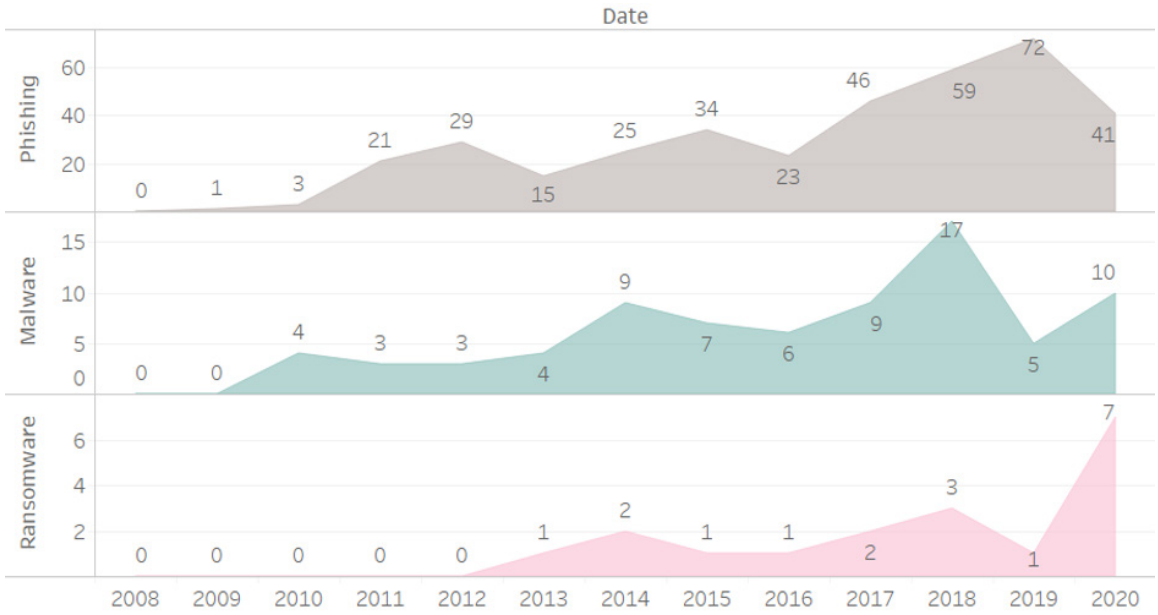

Figure 8. Phishing, Malware, and Ransomware Reporting Timelines

Ransomware attacks have been rapidly increasing in frequency and severity over the past few years, with ransomware incident reports increasing seven-fold year-on-year (Bitdefender, 2020). However, attackers have shifted their focus from developing more sophisticated malware to leveraging social engineering attacks (i.e., phishing) to deliver their malware, ransomware, or simply gain access to an organization's network (Bitdefender, 2020). Threat actors, whether state-sponsored or a cybercriminal group, have realized that exploiting the human tendency to trust is one of the most successful methods of intrusion and acts as a strong delivery method for a range of different attacks (Salahdine & Kaabouch, 2019). Despite consistent numbers surrounding phishing as an attack enabler, ransomware attack types have historically been reported at low levels in the alerts and advisories, especially when compared to the reporting on phishing (Figure 8). Although ransomware continuously maintained low numbers since its first mention within the alerts, 2020 marked a significant spike in ransomware reporting. This indicates that ra-

nsomware operations may have recently begun to specifically target key sectors and critical services, or that the spike in ransomware attacks throughout the past few years has led to an increasing focus on ransomware reporting within the alerts (West & Zentner, 2019).

Hits on fraud ranged from any reporting of fraudulent activities (in a financial context) to the use of ad-fraud attacks. Theft, on the other hand, was not specific to financial theft; it included a range of different types of theft - from identity and data theft to intellectual property or credential theft. Theft acted as a generalized category for this variety due to the low hit counts they would have received individually. Both fraud and theft received significantly low reporting with only occasional and sporadic mentions throughout the past decade (see Appendix C). Additionally, website defacement was a unique TTP compared to the others included in the analysis. Although website defacements have a long history as a prominent attack type, there is only one mention of this TTP in an alert released in 2020 (Appendix C; Maimon et al., 2017). Historically, website defacements may not have been of major concern to critical infrastructure operators. However, the current prevalence of social media platforms, the growing threat of misinformation, and the use of social media enabled spear-phishing likely have a direct impact on a new desire to discuss the threats posed by website defacements (see Maimon et al., 2017).

Limitations

Information sharing has become a vital attribute of cyber defense operations and threat intelligence initiatives (Koepke, 2017; Oosthoek & Doerr, 2020). The results of this study demonstrate a growth in public reporting and represent the U.S. government's increasing willingness to share cyber intelligence. However, various limitations are present within this study's analysis of such reporting. Public reporting on threat activities and vulnerabilities does not reflect the occurrence of cyber-attacks and other underlying events, and the lack of a widely adopted standard for communicating cyber-related information is a significant drawback when attempting to use the same methodology for analyzing other forms of public reporting. Information sharing efforts are crippled by the lack of a widely adopted cybersecurity ontology and reporting mechanism. This creates a challenge in effectively extracting information and driving trend analysis, especially regarding TTPs, as the terminology and frameworks for such an analysis used across the U.S. Government may differ from that of the private sector and academic entities. The integration of a standardized doctrine for cybersecurity terminology and reporting on threat actors, TTPs, targets, and more, would provide a stronger foundation from which future studies can be built.

In addition to working towards a standardized ontology, a stronger framework for extracting TTP information can be integrated into future studies. There are limitations to the use of this study's TTPs dictionary, as it was largely informed by existing literature on adversary TTPs and best judgment based on the terminology used within the alerts and advisories. The TTPs used for this IE provide insight into mentions of some of the most widely known attack types and techniques; however, additional TTP-related terminology likely exists across the sample of alerts and advisories that were not included in this analysis. To tackle this type of challenge, Legoy et al. (2020) drew upon an ML-based information extraction approach by using the MITRE ATT&CK framework to develop a TTP extraction tool that automates the analysis of cyber threat reports. The tool, known as rcATT, enables extraction of TTPs from threat reports by following MITRE's framework and can be used in future studies to further enhance TTP extraction (see Legoy et al., 2020).

The lack of consideration regarding the context of terminology within the alerts when extracting information also served as a prominent limitation. Terms including "communications" and "information technology" (which are defined as sectors of critical infrastructure), may appear in the alerts and yet may not intend to reference an actual sector. For example, "communications" is widely used in discussions surrounding network communications and is not used in relation to the communications sector itself. Incorporating aspects of semantic analysis to better contextualize relevant or irrelevant terminology within the alerts will enhance the performance of an IE and prevent the inclusion of irrelevant terms in the final data outputs (Fedushko & Benova, 2019; Wolff, 2020).

Future Research

The findings of the study present various opportunities for future research. First, there is a significant need for further analysis into how the lack of a widely adopted cybersecurity ontology and reporting mechanism impacts the ability of organizations (whether they be public, private, or academic) to communicate, collaborate, and assess data surrounding cybersecurity risk. The findings also suggest the need for a more comprehensive analysis of open-source information to be conducted to better correlate reporting with underlying cyber events. Private sector reporting on threat actor activities, cyber-attacks, incidents, and even changes in government policy, can build upon this study by driving more substantial correlations between public government reporting and the actual events publicly reported on across each sector of critical infrastructure. It is also worth considering possibilities for how government agencies can enhance public information sharing and the development of standards surrounding cybersecurity reporting. The inclusion of an organization such as the National Institute of Standards and Technology (NIST), which has contributed to the development and wide adoption of various cybersecurity standards and frameworks, could provide the cybersecurity community with a formal approach to tackling the challenges present in public cybersecurity reporting and threat information sharing. Finally, a deeper dive into the specifically attributed adversaries within the alerts can provide an opportunity to assess trends in different threat-actor campaigns while driving insight into the types of threat groups the U.S. Government actively tracks.

Conclusions

This information extraction focused on analyzing three main features of U.S. Government public cybersecurity reporting: sectors of critical infrastructure, the types of adversaries discussed, and the TTPs that are mentioned across the sample of alerts and advisories. As threat actors of all types become increasingly capable of using sophisticated TTPs for their operations, sectors of critical infrastructure remain at high risk of falling victim to a major cyber-attack. The growing threat of cyber-attacks against critical infrastructure calls for an increased emphasis on information sharing across the cybersecurity industry. An analysis of the trends and patterns present within the analyzed alerts and advisories provides insight into the types of cyber-related information provided by U.S. Governmental entities and changes in the risk landscape.

The main findings indicate that reporting on cyber threat activity has seen an exponential increase throughout the past decade, and it is apparent that overall reporting has become more comprehensive and demonstrates a better understanding of the cyber-risk present across different sectors of critical infrastructure. Further analysis of sector-specific reporting generates insight into which sectors of critical

infrastructure are most frequently mentioned in public alerting, enabling a trend analysis that can help depict changes in the threats to critical infrastructure sectors and the possible underlying cyber events that correlate with such changes. Identifying the trends in adversary attribution drives a deeper discussion into the reasons why there has historically been a lack of direct threat attribution, and why only recently has the attribution of such adversaries changed. Understanding which TTPs are consistently reported can also enable CTI analysts to better correlate adversary activities, predict threat behavior, and work to mitigate risks through information sharing and trend analysis. Although the results of this study demonstrate trends in publicly reported threat information, the analysis of such information suggests a significant gap in the way cybersecurity information is communicated between government agencies, private, and academic organizations. Prior research has demonstrated that the lack of a standard approach to public reporting and information sharing causes a disconnect in how cybersecurity information is defined, communicated, and disseminated. These findings drive insight into the way the U.S. Government approaches such reporting while emphasizing the need for the adoption of information sharing standards.

References

Abdyraeva, C. (2020). Cyber Warfare. In *The Use of Cyberspace in the Context of Hybrid Warfare.: Means, Challenges and Trends* (pp. 15–20). OIIP - Austrian Institute for International Affairs. http://www.jstor.org/stable/resrep25102.7

Ahmad, A., Webb, J., Desouza, K., & Boorman, J. (2019). Strategically-motivated advanced persistent threat: Definition, process, tactics and a disinformation model of counterattack. *Computers & Security, 86.* https://doi.org/10.1016/j.cose.2019.07.001

Akhtar, M., Ahmad, Z., & Amin, R. (2020). An efficient mechanism for product data extraction from E-Commerce websites. *Computers, Materials & Continua, 65*(3), 2639–2663. https://doi.org/10.32604/cmc.2020.011485

Antoun, W., Baly, F., & Hajj, H. (2021). AraBERT: Transformer-based model for Arabic language under standing. ArXiv:2003.00104 [Cs]. http://arxiv.org/abs/2003.00104

Bahrami, P. N., Dehghantanha, A., Dargahi, T., Parizi, R. M., Choo, K.-K. R., & Javadi, H. H. S. (2019). Cyber kill chain-based taxonomy of advanced persistent threat actors: Analogy of tactics, techniques, and procedures. *Journal of Information Processing Systems, 15*(4), 865–889. https://doi.org/10.3745/JIPS.03.0126

Bendovschi, A. (2015). Cyber-Attacks – Trends, patterns and security countermeasures. *Procedia Economics and Finance, 28*, 24–31. https://doi.org/10.1016/S2212-5671(15)01077-1

Bitdefender. (2020). Bitdefender mid-year threat landscape report 2020. https://www.bitdefender.com/files/News/CaseStudies/study/366/Bitdefender-Mid-Year-Threat-Landscape-Report-2020.pdf

Brumfield, C. (2019, July 1). What is the CISA? How the new federal agency protects critical infrastructure. *CSO Online*. https://www.csoonline.com/article/3405580/what-is-the-cisa-how-the-new-federal-agency-protects-critical-infrastructure-from-cyber-threats.html

Campbell, R. J. (2018). Electric grid cybersecurity. 31. https://fas.org/sgp/crs/homesec/R45312.pdf

CheckPoint. (2020). Cyber security report 2020. 80. https://pages.checkpoint.com/cyber-security-report-2020.html

Chiticariu, L., Li, Y., & Reiss, F. R. (2013). Rule-based information extraction is dead! Long live rule-based information extraction systems! EMNLP 2013 - 2013 *Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, October, 827–832.

Cimpanu, C. (2020, September 17). First death reported following a ransomware attack on a German hospital. ZDNet. https://www.zdnet.com/article/first-death-reported-following-a-ransomware-attack-on-a-german-hospital/

CISA. (2020, March 24). Critical infrastructure sectors. https://www.cisa.gov/critical-infrastructure-sectors

CISA. (n.d.). Cyber threat source descriptions. Us-Cert.Gov. Retrieved July 4, 2020, from https://us-cert.cisa.gov/ics/content/cyber-threat-source-descriptions

CISA. (2014, June 12). Energy sector | CISA. https://www.cisa.gov/energy-sector

CISA. (n.d.). ICS-CERT advisories. Retrieved October 5, 2020, from https://us-cert.cisa.gov/ics/advisories

Coats, D. R. (2019). National intelligence strategy of the United States of America. https://www.dni.gov/files/ODNI/documents/National_Intelligence_Strategy_2019.pdf

Conti, M., Dehghantanha, A., & Dargahi, T. (2018). Cyber threat intelligence: challenges and opportunities. ArXiv:1808.01162 [Cs], 70, 1–6. https://doi.org/10.1007/978-3-319-73951-9_1

Crowdstrike. (2020). 2020 Global threat report (p. 68). https://www.crowdstrike.com/resources/reports/2020 crowdstrike-global-threat-report/

Dandurand, Luc & Serrano, O.S.. (2013). Towards improved cyber security information sharing. *International Conference on Cyber Conflict*, CYCON. 1-16.

Das, R., & Gündüz, M. Z. (2020). Analysis of cyber-attacks in IoT-based critical infrastructures. *International Journal of Information Security Science, 8*(4), 122–133.

Department of Energy. (2020, May 1). President Trump signs executive order securing the United States bulk-power system. Energy.Gov. https://www.energy.gov/articles/president-trump-signs-executive-order-securing-united-states-bulk-power-system

Department of Justice. (2020, October 19). Six Russian GRU officers charged in connection with worldwide deployment of destructive malware and other disruptive actions in cyberspace. https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and

Dragos. (2020). North American electric cyber threat perspective. 17. https://www.dragos.com/wp-content/uploads/NA-EL-Threat-Perspective-2019.pdf

Dunham, K. (2017, January 19). Tactics, techniques and procedures (ttps) within cyber threat intelligence. Optiv. https://www.optiv.com/explore-optiv-insights/blog/tactics-techniques-and-procedures-ttps-within-cyber-threat-intelligence

E-ISAC. (2016). Analysis of the cyber attack on the Ukrainian power grid. https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf

Edwards, B., Furnas, A., Forrest, S., & Axelrod, R. (2017). Strategic aspects of cyberattack, attribution, and blame. *Proceedings of the National Academy of Sciences, 114*(11), 2825–2830. https://doi.org/10.1073/pnas.1700442114

Eichensehr, K. (2019). The law & politics of cyberattack attribution (SSRN Scholarly Paper ID 3453804). *Social Science Research Network*. https://papers.ssrn.com/abstract=3453804

Eichensehr, K. (2019). Decentralized cyberattack attribution. AJIL Unbound, 113, 213–217. https://doi.org/10.1017/aju.2019.33

Fedushko, S., & Benova, E. (2019). Semantic analysis for information and communication threats detection of online service users. *Procedia Computer Science, 160*, 254–259. https://doi.org/10.1016/j.procs.2019.09.465

Fleury, T., Khurana, H., & Welch, V. (1970). Towards a taxonomy of attacks against energy control systems. 71–85. https://doi.org/10.1007/978-0-387-88523-0_6

G.C., P. S., Deep, R., Raghavendra, V., Doan, A., Sun, C., K., K. G., Zhang, H., Yang, F., Rampalli, N., Prasad, S., Arcaute, E., & Krishnan, G. (2015). Why big data industrial systems need rules and what we can do about it. Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data - SIGMOD '15, 265–276. https://doi.org/10.1145/2723372.2742784

Greenberg, A. (2019). Sandworm: *A new era of cyberwar and the hunt for the Kremlin's most dangerous hackers*. Anchor Books

Geers, K. (2009). The cyber threat to national critical infrastructures: Beyond theory. *Information Security Journal: A Global Perspective, 18*(1), 1–7. https://doi.org/10.1080/19393550802676097

Gritzalis, D., Theocharidou, M., & Stergiopoulos, G. (Eds.). (2019). Critical infrastructure security and resilience: theories, methods, tools and technologies. *Springer International Publishing*. https://doi.org/10.1007/978-3-030-00024-0

Jang-Jaccard, J., & Nepal, S. (2014). A survey of emerging threats in cybersecurity. *Journal of Computer and System Sciences, 80*(5), 973–993. https://doi.org/10.1016/j.jcss.2014.02.005

Karnouskos, S. (2011). Stuxnet worm impact on industrial cyber-physical system security. *IECON Proceed ings* (Industrial Electronics Conference). https://doi.org/10.1109/IECON.2011.6120048

Koehrsen, W. (2018, March 10). Beyond accuracy: Precision and recall. *Medium*. https://towardsdatascience.com/beyond-accuracy-precision-and-recall-3da06bea9f6c

Koepke, P. (2017). Cybersecurity information sharing incentives and barriers. Sloan School of Management at MIT University.

Landi, H. (2020, October 29). Hospitals hit with ransomware attacks as FBI warns of escalating threat to healthcare. *FierceHealthcare*. https://www.fiercehealthcare.com/tech/hospitals-hit-ransomware-attacks-as-fbi-warns-escalating-threat-to-healthcare

Layne, C. (2017). Cyber attacks against critical infrastructure [M.S., Utica College]. http://search.proquest.com/docview/1957428360/abstract/C4EBE4BC74784740PQ/1

Lemay, A., Calvet, J., Menet, F., & Fernandez, J. M. (2018). Survey of publicly available reports on advanced persistent threat actors. *Computers & Security*, 72, 26–59. https://doi.org/10.1016/j.cose.2017.08.005

Levine, N., & Young, M. (2014, August). Leadership trait analysis and threat assessment with profiler plus. In proceedings of ILC 2014 on 8th international Lisp conference (pp. 50-59). https://doi.org/10.1145/2635648.2635657

Lybarger, K. J. (2020). Extracting information from clinical text with limited annotated data [Thesis]. https://digital.lib.washington.edu:443/researchworks/handle/1773/46344

Maglaras, L., Ferrag, M., Derhab, A., Mukherjee, M., Janicke, H., & Rallis, S. (2018). Threats, countermea sures and attribution of cyber-attacks on critical infrastructures. *EAI Endorsed Transactions on Security and Safety, 5*(16).

Maimon, D., Fukuda, A., Hinton, S., Babko-Malaya, O., & Cathey, R. (2017). On the relevance of social media platforms in predicting the volume and patterns of web defacement attacks. 2017 IEEE International Conference on Big Data (Big Data), 4668–4673. https://doi.org/10.1109/BigData.2017.8258513

Mavroeidis, V., & Bromander, S. (2017, September). Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In 2017 European Intelligence and Security Informatics Conference (EISIC) (pp. 91-98). IEEE. http://ieeexplore.ieee.org/document/8240774/

Dion, M., Pacheco, O., McCarthy, J., & Burrow, C. (2009). Cyberpower and Critical Infrastructure Protection: A Critical Assessment of Federal Efforts. MITRE. (n.d.). MITRE ATT&CK®. Retrieved July 29, 2020, from https://attack.mitre.org/#

NCCIC, & FBI. (2016). GRIZZLY STEPPE – Russian malicious cyber activity (Joint Analysis Report JAR-16-20296A). https://us-cert.cisa.gov/sites/default/files/publications/JAR_16-20296A_GRIZZLY%20STEPPE-2016-1229.pdf

Neuendorf, K. A. (2018). Content analysis and thematic analysis. *Advanced Research Methods for Applied Psychology,* 211–223. https://doi.org/10.4324/9781315517971-21

Nussbaum, B. (2019). The growing rumblings of cyberwar. Nature, 280–281.

Olenick, D. (2020, April 28). Israel's water companies suffer cyber-attack. https://www.scmagazineuk.com/article/1681580?utm_source=website&utm_medium=social

Oosthoek, K., & Doerr, C. (2020). Cyber Threat Intelligence: A product without a process?. *International Journal of Intelligence and Counterintelligence,* 1-16.

Ring, T. (2014). Threat intelligence: why people don't share. *Computer Fraud & Security,* 2014(3), 5-9.

Salahdine, F., & Kaabouch, N. (2019). Social engineering attacks: A Survey. *Future Internet, 11*(4), 89. https://doi.org/10.3390/fi11040089

SecureWorks. (2017, May 12). Cyber threat basics, types of threats, intelligence & best practices. https://www.secureworks.com/blog/cyber-threat-basics

Smith, A. (2015). Presidential Policy Directive 21 Implementation: An interagency security committee white paper. 16.

Tableau. (n.d). Data visualization beginner's guide: A definition, examples, and learning resources. Retrieved October 17, 2020, from https://www.tableau.com/learn/articles/data-visualization

Tounsi, W., & Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security*, 72, 212-233.

United States Government Accountability Office (GAO). (2019). Critical Infrastructure Protection: Actions needed to address significant cybersecurity risks facing the electric grid (GAO-19-332; pp. 1–84). https://www.gao.gov/assets/710/701079.pdf

Waltl, B., Bonczek, G., & Matthes, F. (2018). Rule-based information extraction: Advantages, limitations, and perspectives. *Jusletter IT* (02 2018).

Weischedel, R., & Boschee, E. (2018). What can be accomplished with the state of the art in information extraction? A personal view. *Computational Linguistics, 44*, 1–15. https://doi.org/10.1162/coli_a_00331

Weiss, J. (2010). *Protecting industrial control systems from electronic threats. momentum press*. Momentum Press.

West, T., & Zentner, A. (2019). Managing Security Risks: An assessment of u.s. critical cyber infrastructure protection. *SSRN Electronic Journal.* https://doi.org/10.2139/ssrn.3484552

Wolff, R. (2020, August 12). Semantic Analysis: What is it & how does it work? MonkeyLearn Blog. https://monkeylearn.com/blog/semantic-analysis/

Work, J. D. (2020). Evaluating commercial cyber intelligence activity. *International Journal of Intelligence and Counterintelligence, 33*(2), 278-308.

Zetter, K. (2016). *Countdown to zero day: Stuxnet and the launch of the world's first digital weapon.* Crown Publishers.

Appendices
Appendix A

Annotation Guide

The Cyber Alert Information Extraction Scheme (CAIES) is intended to be used to analyze alerts generated by CISA's National Cyber Awareness System (NCAS) and advisories released by ICS-CERT. This information extraction scheme is designed to pull alert information as well as identify the presence of terms that are defined as sectors of critical infrastructure, adversary types, and tactics, techniques, and procedures (TTPs). This annotation guide is designed to:

1. Enable hand coders to evaluate success and generate validation data for the coding scheme.
2. Guide other coders using rule-based information extraction to generate insight into cybersecurity related alerts and publications.
3. Guide how users should properly annotate their data to create useful accuracy, precision, and recall calculations.

What to Code and What Not to Code

1. Codable terms will include any occurrence of TTPs, adversaries, and sectors of critical infrastructure within each alert. As understandings of what is considered/determined to be a TTP, adversary, or designated sector of critical infrastructure is constantly evolving, the most up-to-date understanding of what may be considered applicable to each category should be used and incorporated into both the dictionaries and annotations.

"(CISA) are providing this report to inform the sector about the Dridex malware and variants."	Coded. Malware, a TTP, is coded as it is recognized as a TTP within the TTPsTable dictionary.
"**CRITICAL INFRASTRUCTURE SECTORS**: Energy"	Coded. Energy, a CI-Sector, is coded as a recognized sector of critical infrastructure and is identified as an impacted sector in the alert.
"Specially crafted packets to Port 102/TCP could cause a denial-of-service"	Coded. Denial-of-service is a recognized TTP in the TTPsTable dictionary.
"those users and helps that threat actor move more covertly"	Coded. Threat actor is a recognized adversary type according to the developed ThreatTable dictionary.
"including financial institutions and customers"	Coded. Financial institutions is a recognized sector of critical infrastructure according to the CISectors dictionary.

2. Code TTPs, Adversaries, or CI-Sectors that are recognized in the Alert Title.

"3ve – Major Online Ad Fraud Operation"	Coded. Ad Fraud, a TTP, is mentioned within the title of the alert.
"Potential for Iranian Cyber Response to US Military Strike in Baghdad"	Coded. Iranian Cyber Response, which is recognized as a Nation-State adversary appears in the title.
"Ransomware Impacting Pipeline Operations"	Coded. Ransomware, a TTP, is mentioned within the title of the alert.

3. Headlines (NOT alert titles) that mention TTPs, adversaries, or targets are to be coded.

"Dridex-related Phishing Attributes"	Coded. Phishing, a TTP, is a recognized TTP and is included in a headline.
"APT TTPs and Corresponding Mitigations"	Coded. APT, is coded as it is recognized in the ThreatTable dictionary.
"Cyber Operations Publicly Attributed to DPRK by U.S. Government"	Coded. DPRK, a Nation-State adversary, is a defined adversary and is included in a headline.

4. Do not code links.

"http://blog.talosintelligence.com/2017/06/worldwide-ransomware-variant.html"	Not Coded. Ransomware, a TTP, is contained within a link.
"https://www.justice.gov/opa/pr/justice-department-announces-actions-disrupt-advanced-persistent-threat-28-botnet"	Not Coded. Advanced Persistent Threat and Botnet, an Adversary and a TTP, is contained within a link.
"https://www.alertlogic.com/resources/threat-reports/dridex-malware-has-evolved-to-locky-ransomware"	Not Coded. Malware and ransomware, two TTPs, are contained within a link.

"https://www.us-cert.gov/report-phishing"	Not Coded. Phishing, a TTP, is contained within a link.
"https://www.energy.gov/sites/prod/files/2014/03/f13/ONG-C2M2-v1-1_cor.pdf"	Not Coded. Energy, a CI-Sector is contained within a link.

5. Do not code terms if they are not in reference to the term's intended category

"the IEC 61850-MMS communication needs to be activated"	Not Coded. Communications is not used in reference to the communications sector.
"below—inspected inbound financial request messages for specific primary account numbers"	Not Coded. Financial is not used in reference to the financial sector.
"modification of firewall rules to facilitate peer-to-peer communication for extraction of data"	Not Coded. Communication is not used in reference to the communications sector.

How to Hand Code

Hand coding involves placing tags into documents. The tags consist of a code with a tilde on each side, such as ~TTP~ or ~CI-Sector~. Tags cannot contain spaces. The type and placement of tags will vary for each coding scheme. There are three types of codes used for CAIES hand-coding.

1. TTP codes (~TTP~) to identify the tactics, techniques, and procedures mentioned within an alert.
2. Adversary codes (~Adversary~) to identify the adversaries mentioned within an alert.
3. Critical Infrastructure Sector codes (~CI-Sector~) to identify the sectors mentioned within an alert.

Each of the identified tag types must be placed in front of the appropriate term within the sentence.

"New York, indicted nine ~adversaries~ Iranian nationals"
"Using ~TTP~ social engineering tactics to perform online research"
"CRITICAL INFRASTRUCTURE SECTORS:** ~CI-Sector~ Chemical, ~CI-Sector~ Critical Manufacturing, ~CI-Sector~ Energy"
"unauthenticated ~Adversary~ attacker may be able to compromise a vulnerable VPN"
"devices—coupled with a ~Adversary~ Russian government campaign

The code cannot be placed after the term as it will recognize the value following the enumeration as the coded item.

"This could deceive users or thwart malware ~TTP~ detection methods"
"An attacker ~Adversary~ could inject arbitrary JavaScript in a specially crafted URL"
"law enforcement efforts in protecting the financial sector ~CI-Sector~.
"This vulnerability allows remote attackers ~Adversary~ cause a denial-of-service ~TTP~ condition due to a lack of proper validation

If a defined term contains two words or more, insert the code prior to the first word in the phrase.

"~Adversary~ Russian State-Sponsored Cyber Actors Targeting Network Infrastructure Devices"
"~Adversary~ Malicious cyber actors can attack and compromise these unsecure systems"
"A ~TTP~ remote code execution vulnerability exists in Windows Remote
Desktop Gateway"
"~CI-Sector~ Water and Wastewater Systems"

If a sentence contains multiple defined terms, all relevant terms must be coded.

"Inform and educate employees on the appearance of ~TTP~ phishing messages, especially those used by the ~Adversary~ hackers for distribution of ~TTP~ malware in the past."
"Successful exploitation of these vulnerabilities may allow ~TTP~ information
disclosure, ~TTP~ privilege escalation, or ~TTP~ remote code execution."
"This vulnerability allows remote ~Adversary~ attackers to ~TTP~ execute arbitrary code due to the lack of proper"
"CRITICAL INFRASTRUCTURE SECTORS:** ~CI-Sector~ Critical Manufacturing, ~CI-Sector~ Energy, ~CI-Sector~ Water and Wastewater Systems"
"Successful exploitation of these vulnerabilities may allow ~TTP~ arbitrary code
execution, a partial ~TTP~ denial-of-service condition, or ~TTP~ information disclosure."
"Successful exploitation of this vulnerability could allow a remote ~Adversary~ attacker to gain ~TTP~ elevated privileges for ~TTP~ remote code execution."

Appendix B

Breakdown of Sector-Specific Reporting by Year
(note: scales differ among graphs)

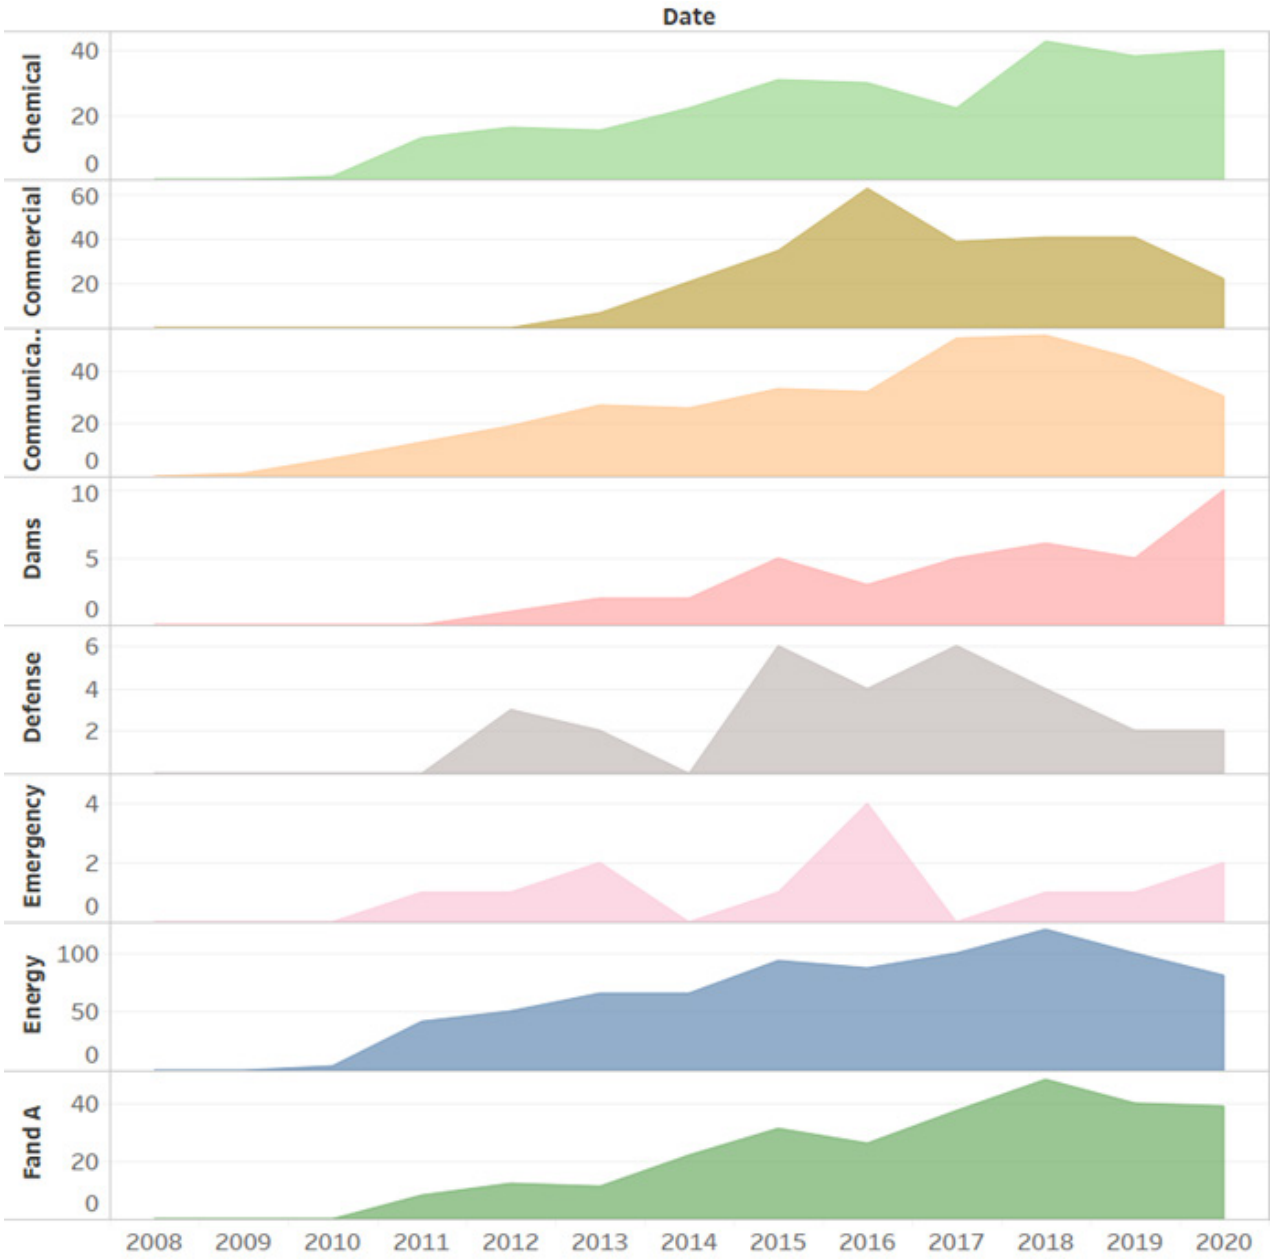

Appendix B

Breakdown of Sector-Specific Reporting by Year
(note: scales differ among graphs)

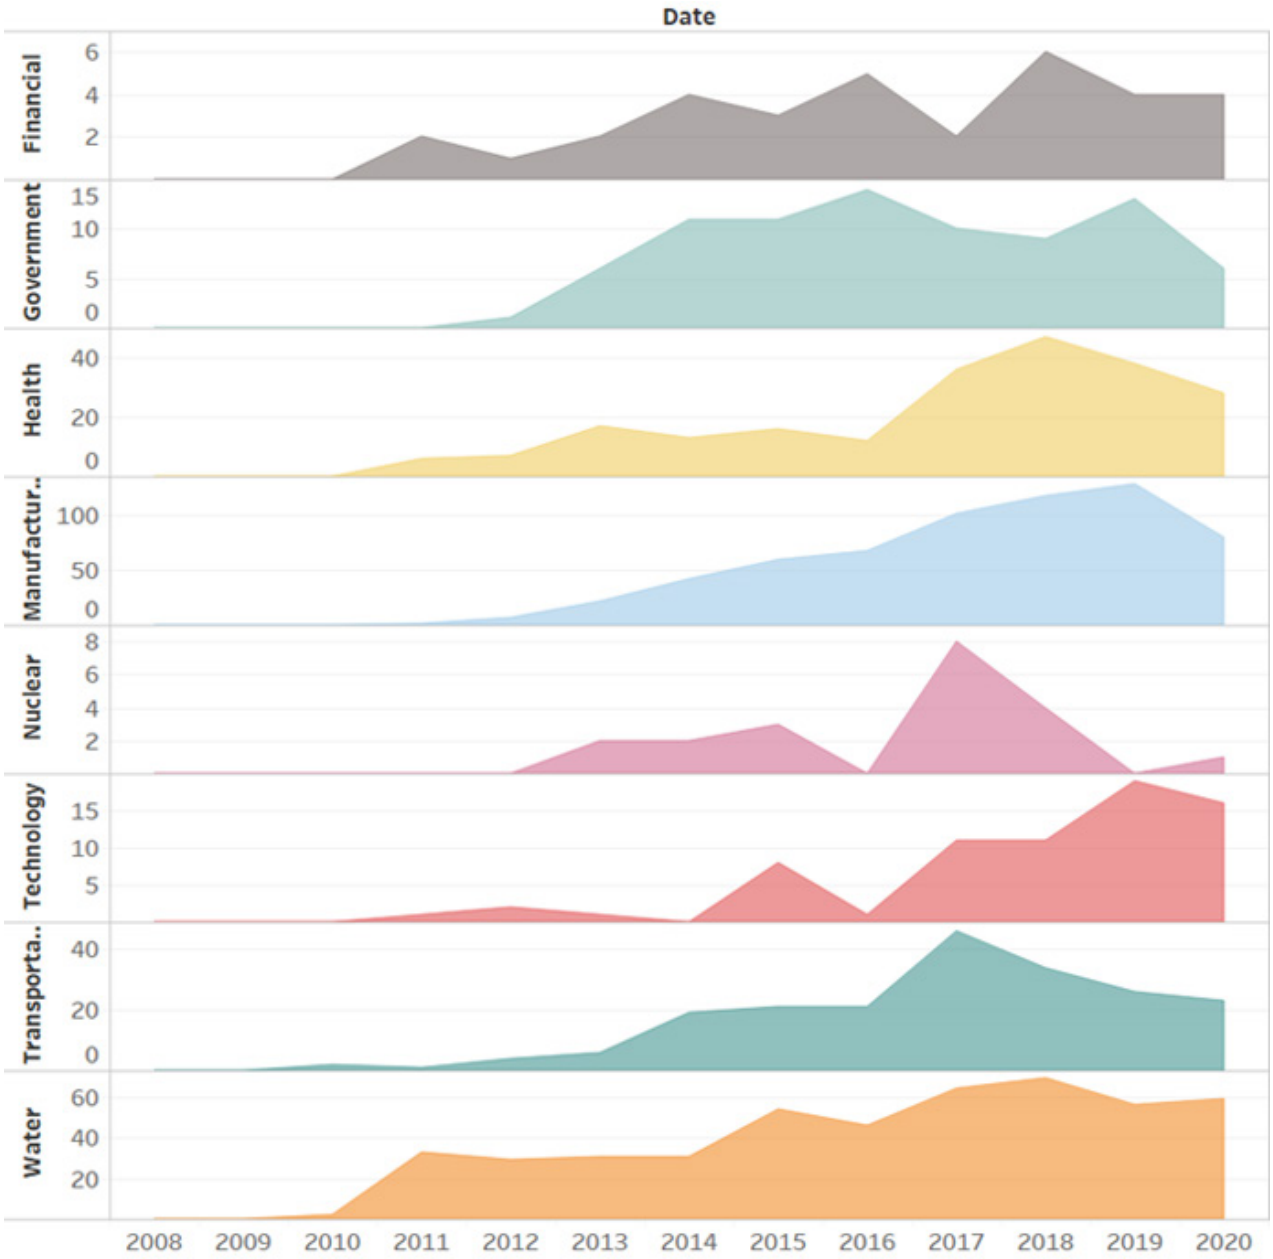

Appendix B: Breakdown of Sector-Specific Reporting by Year

Appendix C

Breakdown of Sector-Specific Reporting by Year
(note: scales differ among graphs)

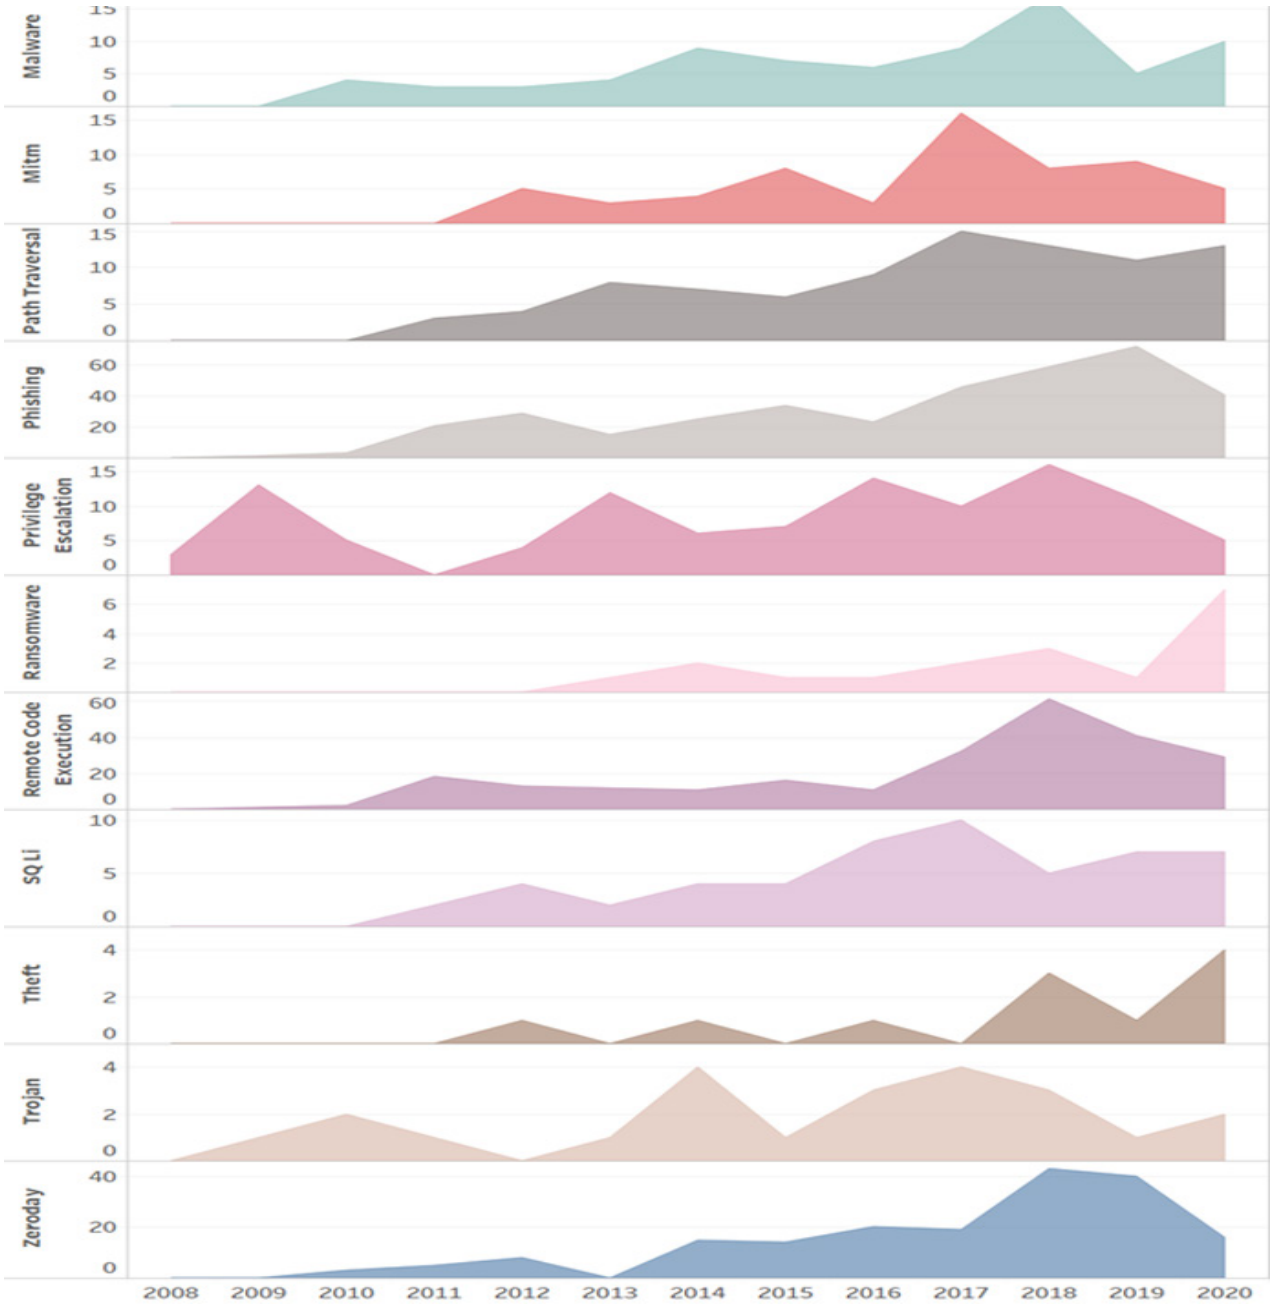

Appendix C

Breakdown of Sector-Specific Reporting by Year
(note: scales differ among graphs)

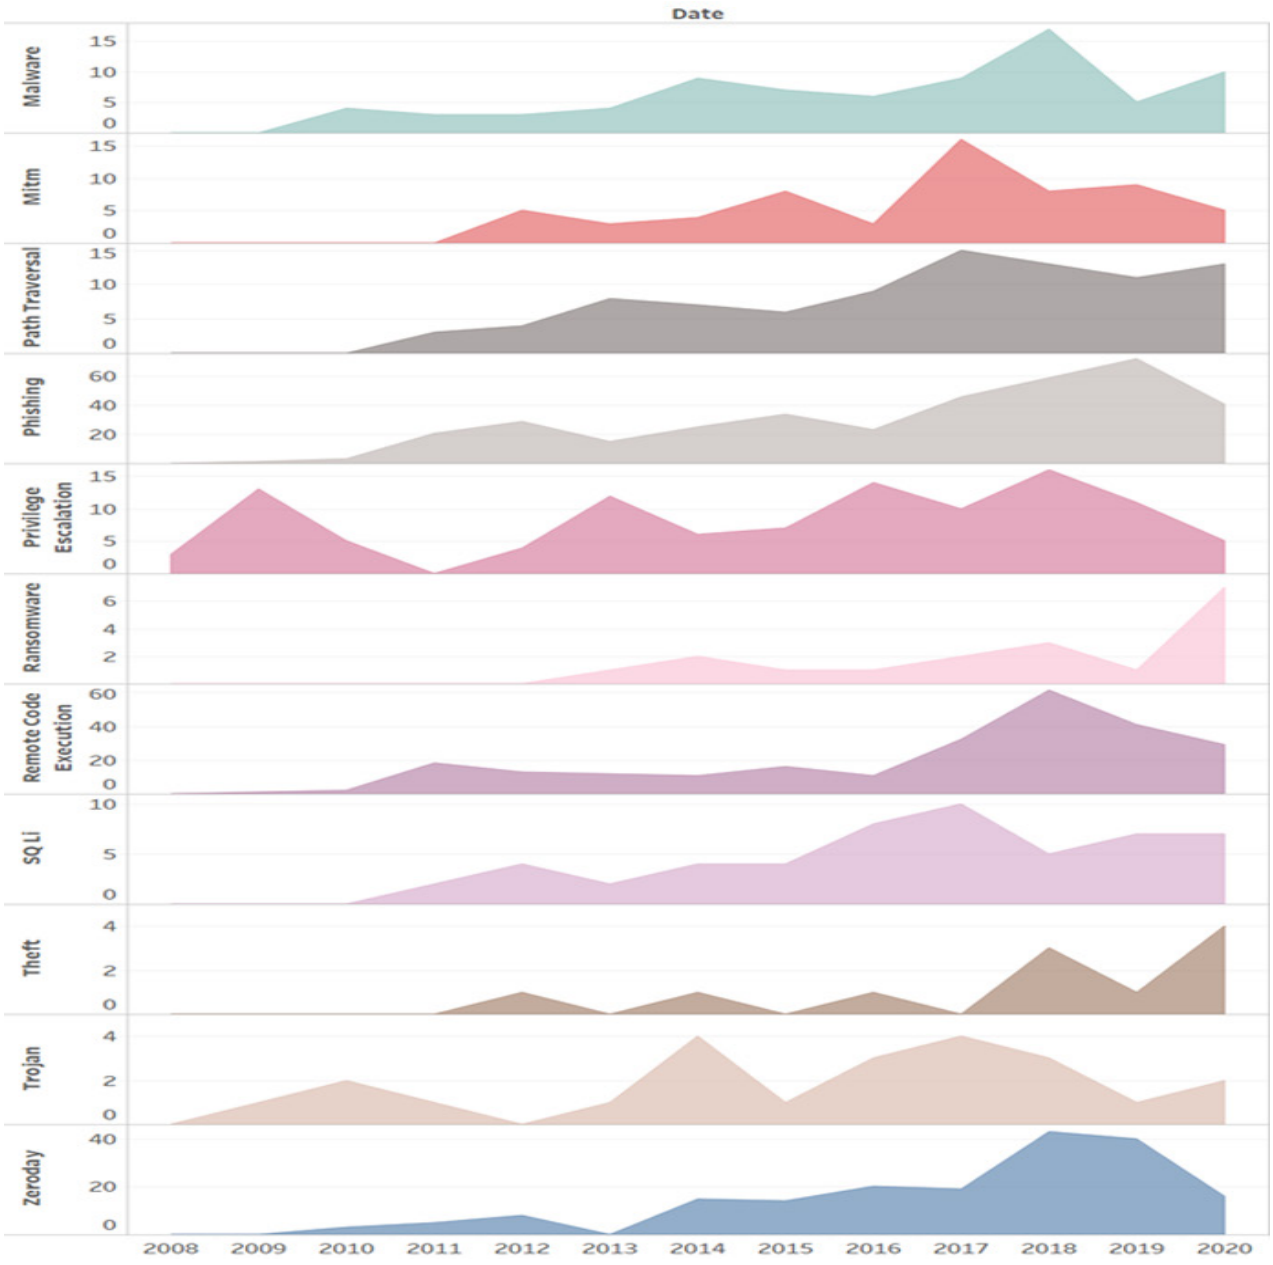

Appendix C: Breakdown of TTP Specific Reporting by Year