Assessing Internet-wide Cyber Situational Awareness of Critical Sectors

Conference Paper · August 2018

DOI: 10.1145/3230833.3230837

5 authors, including:

Martin Husák
Masaryk University
27 PUBLICATIONS **132** CITATIONS
SEE PROFILE

Nataliia Neshenko
Florida Atlantic University
7 PUBLICATIONS **35** CITATIONS
SEE PROFILE

Morteza Safaei Pour
University of Texas at San Antonio
6 PUBLICATIONS **16** CITATIONS
SEE PROFILE

Elias Bou-Harb
University of Texas at San Antonio
73 PUBLICATIONS **554** CITATIONS
SEE PROFILE

Assessing Internet-wide Cyber Situational Awareness of Critical Sectors

Martin Husák
Institute of Computer Science
Masaryk University
Brno, Czech Republic
husakm@ics.muni.cz

Nataliia Neshenko
Cyber Threat Intelligence Laboratory
Florida Atlantic University
Boca Raton, USA
nneshenko2016@fau.edu

Morteza Safaei Pour
Cyber Threat Intelligence Laboratory
Florida Atlantic University
Boca Raton, USA
msafaeipour2017@fau.edu

Elias Bou-Harb
Cyber Threat Intelligence Laboratory
Florida Atlantic University
Boca Raton, USA
ebouharb@fau.edu

Pavel Čeleda
Institute of Computer Science
Masaryk University
Brno, Czech Republic
celeda@ics.muni.cz

ABSTRACT

In this short paper, we take a first step towards empirically assessing Internet-wide malicious activities generated from and targeted towards Internet-scale business sectors (i.e., financial, health, education, etc.) and critical infrastructure (i.e., utilities, manufacturing, government, etc.). Facilitated by an innovative and a collaborative large-scale effort, we have conducted discussions with numerous Internet entities to obtain rare and private information related to allocated IP blocks pertaining to the aforementioned sectors and critical infrastructure. To this end, we employ such information to attribute Internet-scale maliciousness to such sectors and realms, in an attempt to provide an in-depth analysis of the global cyber situational posture. We draw upon close to 16.8 TB of darknet data to infer probing activities (typically generated by malicious/infected hosts) and DDoS backscatter, from which we distill IP addresses of victims. By executing week-long measurements, we observed an alarming number of more than 11,000 probing machines and 300 DDoS attack victims hosted by critical sectors. We also generate rare insights related to the maliciousness of various business sectors, including financial, which typically do not report their hosted and targeted illicit activities for reputation-preservation purposes. While we treat the obtained results with strict confidence due to obvious sensitivity reasons, we postulate that such generated cyber threat intelligence could be shared with sector/critical infrastructure operators, backbone networks and Internet service providers to contribute to the overall threat remediation objective.

CCS CONCEPTS

• **Security and privacy → Network security**; *Intrusion/anomaly detection and malware mitigation*; *Denial-of-service attacks*;

KEYWORDS

network security, network scanning, DDoS, critical infrastructure

ACM Reference Format:
Martin Husák, Nataliia Neshenko, Morteza Safaei Pour, Elias Bou-Harb, and Pavel Čeleda. 2018. Assessing Internet-wide Cyber Situational Awareness of Critical Sectors. In *Proceedings of ARES '18.* ACM, New York, NY, USA, 7 pages. https://doi.org/00.0000/0000000.0000000

ARES '18, 2018, Hamburg, Germany

ACM ISBN 000-0-0000-0000-0/00/00…$15.00
https://doi.org/00.0000/0000000.0000000

1 INTRODUCTION

In recent years, we have witnessed the large-scale adoption of Internet connectivity in various sectors of industry and commerce. The rise of paradigms such as the Internet of Things (IoT) and Cyber-Physical Systems (CPS) illustrates that not only end users, but almost any device/critical infrastructure is now connected to the Internet. Consequently, a number of threats have abused such connectivity nature and have been historically reported. For instance, targeted malware that aimed at infecting critical infrastructure include the notorious `Stuxnet` [7], `Havex` and `Industroyer` [21]. Along the same line of thought, contemporary malware such as `Mirai`, `Persirai`, `Hajime` and `BrickerBot` have been exploiting IoT infrastructure to launch devastating attacks, threatening the resiliency of the Internet at large. Indeed, even a commodity infected device, such as a desktop, may cause significant harm when deployed and operated in a critical infrastructure's network, as observed with the recent `WannaCry` ransomware as it targeted worldwide health sectors [11].

Having noted the above, it is quite understandable that there exists an utmost need (from researchers and practitioners) for cyber threat intelligence and wide-area cyber situational awareness addressing critical infrastructures [1]. Nevertheless, it is known that it is very laborious (or even infeasible) to have access to traces of network traffic or cyber attacks from various industries, not to mention the limited coverage of such data (e.g., covering only customers of Kaspersky [19]), in addition to the unwillingness of certain sectors to share such information [29]. Given such hindering challenges, in this work, we explore unique and complementary empirical data in an attempt to shed light on Internet-scale maliciousness of critical sectors. However, the attribution of IP addresses with a business/critical sector is non-trivial. To this end, we are

involved with a large-scale collaborative effort to obtain private and rare information related to allocated/operated IP blocks per Internet business/industry; we hope that we can discuss the details of this effort at the conference. Thus, we employ such information to attribute Internet-scale maliciousness to various sectors. By generating such threat intelligence, we are capable of assessing the cyber security posture of Internet-wide sectors, with a special focus on critical infrastructures. We postulate that such information could be distributed to the operators of such entities, computer emergency response teams and Internet service providers to aid in the global remediation objective.

In this context, we scope the contributions of this work by posing the following research questions:

(1) Given the lack of empirical data that can be analyzed from within various sectors, including critical infrastructure, in addition to the complementary logistics and privacy issues, how can one assess the Internet-scale cyber security posture of such sectors?
(2) What insights and inferences can one generate by analyzing and characterizing sector-related empirical data, which could be used for effective cyber threat intelligence?

To answer the aforementioned research questions, we investigate macroscopic data collected by darknets, which represent assigned, routable yet unused IP address spaces, where mostly network scans, backscatter, and misconfigured network traffic can be observed. We identify network scanning activities, which typically indicate an infected host (after filtering benign scanning activities from known entities), and DoS attacks, for which we can infer victims' IP addresses. Consequently, we attribute IP addresses of scanning hosts and DDoS victims hosted at or targeting various business/critical sectors. A metric of scan to DDoS ratio is also introduced to provide further insights related to the cyber security posture of such realms.

This paper is organized as follows. Section 2 presents the related work. Section 3 details our approach related to data collection and analysis. The results of the executed empirical measurements are presented and discussed in Section 4. Section 5 concludes the paper and pinpoints few topics, which pave the way for future work.

2 RELATED WORK

Two main approaches can be used to identify and characterize Internet-facing devices hosted in critical sectors, namely, active and passive measurements.

In the area of active measurements for device characterization, Cui and Stolfo [9] executed a large-scale active probing of the Internet space to uncover close to half a million vulnerable embedded devices. Further efforts backed by Internet-wide scanning led to the development of specialized tools and databases, where findings of such scans are stored. Databases backed by Internet-wide scanning, such as Shodan [22] and Censys [10], can be used to search for specific devices facing the Internet, including IoT and critical infrastructure. Bodenheim et al. [4] evaluated Shodan's ability to scan and index Industrial Control Systems (ICS).

An alternative approach to characterizing network traffic, with a special focus on cyber security, is based on passive measurements. Internet background radiation [28], including network scanning [5], backscatter [2], and misconfigured traffic, is a significant source of cyber threat intelligence. Darknets (or network telescopes), as noted earlier, are routable, yet unpopulated segments of IP address space specifically deployed to capture Internet background radiation. Any network traffic targeting the darknet is, by nature, suspicious and unsolicited. This includes network scans, that are typically performed by infected devices trying to spread malware (or search for vulnerabilities), and backscatter, which is an accompanying phenomenon of many DDoS attacks, studied by Balkanli and Zincir-Heywood [2], Blenn et al. [3], and many others. Galluscio et al. [17] were the first to empirically evaluate Internet-scale exploitations of IoT devices using data from darknets and correlating them with data from Shodan. In a similar work, Fachkha et al. [15] conducted passive measurements to analyze attackers' intentions when targeting protocols of Internet-facing CPS.

Apart from darknets, honeypots were also used in the analysis of critical infrastructure exploitations and threat intelligence. For example, Conpot [24] is a well-known ICS/SCADA honeypot that is quite popular among researchers and practitioners. Conpot is able to simulate various types of ICS and SCADA devices, that attackers may interact with. A mobile ICS honeypot dubbed as HosTaGe, as demonstrated by Vasilomanolakis et al. [27], is capable of emulating nuclear power plants, water distribution plants, etc. It is not unusual to set up a large and detailed decoy infrastructure to attract attackers and generate attack signatures [26]. Similar research and development efforts were conducted in the IoT realm. Honeypots, such as IoTPOT [23] were developed and deployed to analyze IoT-specific threats and attacks. IoTPOT emulates Telnet services of various IoT devices running on different CPU architectures, which allows capturing and analyzing various types of malware samples targeting IoT. Guarnizo et al. [18] presented the Scalable high-Interaction Honeypot (SIPHON) platform for IoT devices. The authors were able to mimic various IoT devices connected to the Internet to attract malicious network traffic, leveraging worldwide wormholes and a few physical devices. Characterization of the malicious network traffic, including employed protocols, is also provided.

In this work, we employ a significant amount of macroscopic darknet data to infer Internet-scale maliciousness, including activities generated from machines hosted by numerous business sectors that typically would not be willing to share any information about their networks or security breaches. Further, we employ the obtained private data related to sector-allocated IP blocks to attribute such misdemeanors to critical sectors.

3 PROPOSED APPROACH

In this section, we provide details on our approach for inferring Internet-wide malicious activities related to critical sectors using passive measurement data from network telescopes.

3.1 Collecting Darknet Data

Having an insight into the network traffic of various sectors is challenging. It is relatively infeasible to get visibility into the operational networks of various sectors, especially those of critical nature. The issue is not only technical, but such activities would also require the approval of the network owners, who are in most cases unwilling to disclose anything about their networks. Nevertheless, large-scale

assessment of malicious activities can be performed even without direct access to the networks of our interest. To this end, passive measurements and analysis rendered by analyzing darknet data is quite an effective methodology to achieve the latter objective; a darknet [16] represents a partial view of the entire Internet address space and often consists of numerous network sensors distributed throughout the Internet.

For the purpose of this work, we leverage real-time darknet data from a /8 network telescope through our collaboration with the Center for Applied Internet Data Analysis (CAIDA) [6]. Such a large darknet represents 1/256 of the total IP address range and, thus, provides a very thorough vantage point into Internet-wide unsolicited network traffic, namely network scanning, DDoS backscatter, and traffic from misconfigured devices.

To infer network scanning and DDoS backscatter, we exploit flow-based parameters of anomalous network traffic. A darknet flow is defined as a series of consecutive packets from the same source IP address and sharing other selected parameters, such as IP protocol, port numbers, and TCP flags. The number of packets per flow is calculated within a certain time window. If a predefined threshold of packets is exceeded, the flow is marked as either a scan or a backscatter, depending on additional criteria. For example, a flow of packets with TCP SYN flag, the same destination port, and various destination IP addresses suggests horizontal network scanning, while a flow of packets with TCP SYN ACK flags and the same source port indicates a DDoS backscatter from a TCP SYN flood. Other characteristics are used to detect vertical and strobe scans, and ICMP backscatter from UDP floods. In this work, we set the threshold to 64 packets borrowed from [20].

3.2 Identifying Critical Sectors

Once a set of suspicious IP addresses is identified by exploring darknet data, we proceed with attributing the IP addresses to their corresponding sectors/industry. Manual attribution would include DNS and WHOIS querying, active probing, and interpretation of the results, which would be too laborious and time-consuming. To this end, as noted earlier, we are involved in a collaborative effort to access and collect private information related to Internet-accessible IP blocks. Using these information, we attribute the unsolicited and inferred IP addresses with their hosting sectors.

Attributing IP addresses with commercial and industrial sectors is only the first step in assessing the cyber security posture of critical infrastructures. Indeed, there is a need to identify the sectors that would be deemed as being critical. While this might appear trivial, it is indeed challenging given that there is no commonly accepted definition of Critical Infrastructures (CI) or Critical Information Infrastructures (CII). For example, the United States Department of Homeland Security (DHS) defines 16 sectors of industry that together fold the CI sector [25]. The full list of critical infrastructure is listed in Table 1.

Alternatively, the European Union (EU) perceives critical infrastructure as an asset or system which is essential for the maintenance of vital societal functions [12]. The list of critical infrastructures can be found in EU directives, such as the one provided by the Council Directive 2008/114/EC [8]. The directive defines two sectors, namely, the European Critical Infrastructure (ECI), and Energy and Transport. The Energy sector covers Electricity, Oil, and Gas sub-sectors while the Transport sector covers Road, Rail, Air, and Inland waterway transport, as well as Ocean and short-sea shipping and ports.

Table 1: Critical sectors as listed by DHS [25]

Chemical	Financial Services
Commercial Facilities	Food and Agriculture
Communications	Government Facilities
Critical Manufacturing	Healthcare and Public Health
Dams	Information Technology
Defense Industrial Base	Nuclear Reactors, Materials, and Waste
Emergency Services	Transportation Systems
Energy	Water and Wastewater Systems

Although we can find many examples of critical infrastructures in the physical world following the above-mentioned definitions, the attribution might not be as clear in cyberspace. Thus, at least in the EU, the term CII has appeared to describe critical infrastructure from the perspective of cyberspace and cyber security. Readers are kindly referred to related documents by ENISA [13, 14] for more information on the topic and we hope that we can discuss the notion of CII attribution and labeling in the conference. To proceed with data analysis, we carefully select those (critical) sectors, which appear in either lists provided by DHS and EU.

3.3 Data Analysis

To analyze the data, apart from scanning and DDoS backscatter inference, we introduce a metric to assess the differences in malicious activities regarding the concerned sectors. The metric defines the ratio of network scanning to DDoS attacks related to a given sector, as perceived and computed from the share of a given sector's scans and DDoS attacks. Network scanning from a hosting sector indicates infected hosts (and thus represent its global security posture/health), while targets of higher interest would presumably attract more DDoS attacks. Thus, a below average scan to DDoS ratio would suggest a sector with more infected hosts of less significance, while an above average ratio would suggest a sector with hosts that are more prone to DDoS than to infections.

Using such metric, we are interested in the differences between individual sectors, as well as between critical and other sectors. Assuming the critical sectors would be more secured and well maintained, there should be far less scans originating from such networks. On the other hand, critical sectors host many systems that are tempting targets for DDoS attacks. In contrary, end-user networks, such as networks of Internet service providers, are assumed to be more likely to be infected with malware, but not being interesting for attracting DDoS attacks. Setting metrics for such behavior would aid in characterizing such and other networks (for which we might currently not know the corresponding sector because of lack of corresponding available IP block data).

4 EMPIRICAL EVALUATION

In this section, we employ the approach proposed in the previous section to elaborate on the results of passive measurements and

Table 2: Events and distinct IP addresses per day.

Day	2018-02-19	2018-02-20	2018-02-21	2018-02-22	2018-02-23	2018-02-24	2018-02-25
Network scanning	8,419,918	8,064,067	8,512,733	7,959,108	7,983,071	7,935,752	7,975,506
Distinct scanning IPs	1,795,788	1,787,678	1,896,957	1,780,817	1,784,485	1,767,198	1,787,251
DDoS attacks	52,030	29,404	184,921	32,711	32,285	29,581	29,437
Distinct DDoS victim IPs	6,782	4,810	92,290	8,058	7,160	5,190	5,054

analytics. The dataset is described first, followed by the results of (critical) sector attribution and data analysis.

4.1 Collected Data

For the empirical evaluation, we employed a dataset of approximately 16.8 TB of darknet data collected during one week at the end of February 2018. Using the approach described in Section 3, we were able to infer 57,240,254 events in total, around 8 million events per day. Out of this number, a vast majority of entries were network scanning events. The numbers of DDoS attacks observed via backscatter were significantly lower. Table 2 provides a breakdown of event types per day.

Please note that the number of distinct IP addresses is typically several times lower than the number of inferred events. The IP addresses represent either scanning hosts or victims of DDoS attacks. The reason behind this is that the events are repeated multiple times per day or week at different times, which is not reflected by the inference algorithm. However, the observation of such repetitions provides additional information on scanners and DDoS attacks. For instance, they might indicate aggressive or persistent scanners or long-term DDoS campaigns against a specific victim. As we can note in Table 2, only slightly above 20 % of events per day contains a distinct IP address. Further, numerous IP addresses were observed repeatedly in the consecutive days or during the whole week. For the purpose of further analysis, we only focused on distinct IP addresses.

4.2 Critical Sector Attribution

Having inferred network scanning and DDoS attacks from darknet data, we proceeded to attribute hosting sectors to scanning IP addresses and DDoS victims.

The first task is to attribute the events and their corresponding IP addresses with their hosting sector. Out of the 57,240,254 events related to 8,101,292 distinct IP addresses observed during the measurement period, we were able to attribute a hosting sector to 92.08 % of distinct IP addresses related to 86.73 % of the events. Specifically, we observed 7,988,706 distinct scanning IP addresses and 112,586 distinct targets of DDoS attacks. However, there is a discrepancy between scanning IP addresses and DDoS targets attributed with a sector. While only 13.14 % of scans were originated from 7.72 % unattributed IP addresses, we were not able to retrieve a hosting sector for victims of 31.70 % DDoS attacks against 22.51 % of distinct targets.

As illustrated in Figure 1, the vast majority of inferred events were attributed to the Telecommunications and Internet service providers' sectors. In addition, the third most frequently attributed sector was that of Internet hosting services. These three sectors pose a unique position on the Internet, and it is only natural that they cover so many IP addresses, such as home Internet connections and mobile devices.

The second task is to identify critical sectors and to filter the events (scanning and DDoS attacks) associated with the given sectors. However, in contrary to the previous task, we cannot automate this process as, to the best of our knowledge, there are no available machine-readable lists of critical sectors. We had to scrutinize the lists provided by DHS [25] and EU [8] presented in Section 3 and correlate these lists with the list of sectors assigned to IP address blocks (which is available to us). The attribution of critical sectors required numerous educated guesses due to several factors. The list provided by DHS actually covers almost every sector of industry and commerce; criticality often depends on additional parameters, such as the number of users and/or people involved. For example, DHS would consider a large hotel or a shopping center as a critical infrastructure, as it attracts a lot of people. Further, it is not clear which manufacturing is critical because we cannot distinguish between manufacturing types. On the contrary, the list provided by EU is more narrow and can be directly mapped to the Transportation and Utilities sectors.

In total, we observed 49 distinct sectors, out of which we identified 6 sectors as critical. The share of critical sectors in the context of the inferred security events is, according to our findings, less than 1 %. To be exact, critical sectors hosted 0.14 % scanning hosts and 0.31 % of DDoS victims. A graphical breakdown, including the critical sectors, is depicted in Figure 1. For the sake of simplicity, we merged similar sectors, mostly non-critical and less frequent ones. Notable exceptions are the Health sector, which consists of sectors of Health and Hospitals, and the Financial sector, which consists of sectors related to Banking, Finances, Insurance, and Real Estate. Further, the Motor Vehicles sector was merged with the Transportation sector. In addition, the governmental sector appeared under multiple labels, depending on the government level, i.e., municipal, county, state, federal, and general. All the government levels were merged into one sector.

4.3 Scan to DDoS Ratio

Herein, we present the results of data analysis of network scans and DDoS attacks associated with various sectors of industry and commerce. Table 3 shows the scan to DDoS ratios for selected highly-frequented sectors. This metric allows us to see which sectors are more prone to host infected devices or to host services that attract DDoS attacks. Two representative examples may be found in the sectors related to Telecommunications and Internet hosting services. Telecommunications host large numbers of end-user devices, which might be infected at large-scale, but are not interesting targets of DDoS attacks. Therefore, the ratio is above average and is highlighted in gray in Table 3. In contrary, Internet hosting services

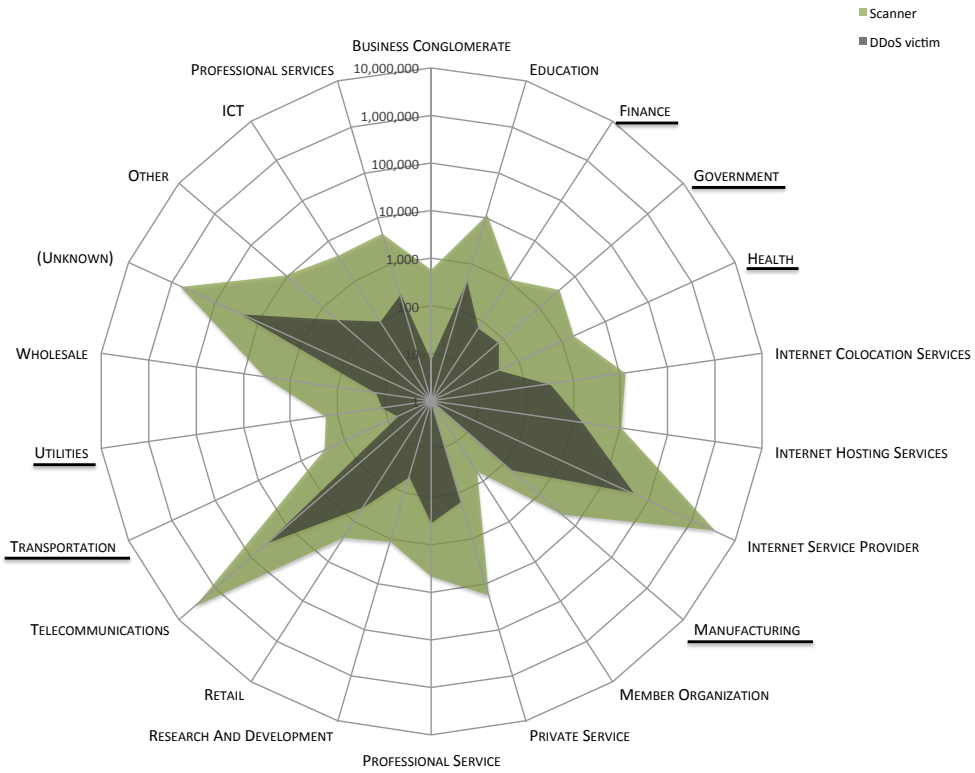

Figure 1: Scanners and DDoS victims per sector; critical sectors are underlined.

host a number of services that are likely to be targeted by DDoS attacks, while some level of security is assumed, which would imply a lower number of infected hosts and network scans. Therefore, the ratio is below average and is highlighted in green in Table 3. The sectors with around-average (± 25 %) ratios are also summarized in Table 3 but are not highlighted.

Critical sectors are assumed to follow characteristics similar to Internet hosting services and, as we can see from the results, they really do. We can infer that most clearly in the case of the Financial sector, which includes banks, and also in sectors such as Manufacturing and Utilities. Further, we can also note this characteristic when we combine all the critical sectors in one group.

However, Government, Health, and Transportation sectors have only an around-average ratio, which would indicate worse situation than, e.g., in the Financial sector. This is especially alarming in the Health sector, that includes hospitals. On the other hand, no critical sector have the ratio significantly higher than average.

An interesting observation is the somehow low ratio for events associated with IP address ranges with an unknown sector. This could indicate a presence of IP address ranges of sectors that are well maintained and/or prone to DDoS attacks, possibly even critical sectors. Regarding this information, further investigation into these unknown IP ranges will be executed in the near future.

Table 3: Scan to DDoS share ratio of selected sectors.

Top-10 non-critical sectors			
Sector	**Scans (%)**	**DDoS (%)**	**Ratio**
Telecommunications	47.668	33.049	1.442
Internet Service Provider	43.404	40.583	1.069
(unknown)	7.717	22.505	0.343
Private Service	0.224	0.134	1.671
Internet Colocation Services	0.157	0.292	0.538
Education	0.154	0.388	0.397
Internet Hosting Services	0.135	1.351	0.100
Other	0.137	0.341	0.402
Professional Service	0.059	0.314	0.187
ICT	0.053	0.085	0.623
Critical sectors			
Sector	**Scans (%)**	**DDoS (%)**	**Ratio**
Manufacturing	0.053	0.139	0.383
Government	0.044	0.064	0.693
Health	0.024	0.032	0.736
Finance	0.014	0.056	0.247
Transportation	0.004	0.005	0.684
Utilities	0.002	0.010	0.219
All critical sectors combined	0.140	0.306	0.460
Average ratio (all sectors)			**0.681**

5 CONCLUSION AND FUTURE WORK

In this short paper, we employ macroscopic darknet data to infer Internet-scale network scanning and DDoS events. We attribute IP addresses of scanners and DDoS victims with their corresponding hosting sectors, including critical sectors, to assess which sectors are prone to infections and/or DDoS attacks. In total, we were able to assign 49 distinct sectors to more than 92 % of IP addresses associated with scanning or DDoS activities. Out of those sectors, 6 were identified as critical, representing around 0.14 % of distinct IP addresses. While such results represent only a small percentage, they still indeed include over 11,000 distinct probing IP addresses and over 300 distinct victims of DDoS attacks in only one week. Our results show a potential metric for characterizing (critical) sectors based on a ratio of inferred network scanning events to DDoS attacks. Thus, we can observe which sector is more likely to be infected and which is more tempting for DDoS attacks. It is worth noting that our approach involves security events associated with sectors that are usually not willing to disseminate information about cyber attacks targeting them, e.g., due to brand damage as it is common in the banking sector. Further, our approach is Internet-wide and not limited to customers of a certain DDoS protection service.

Indeed, this work renders a first attempt to execute Internet-wide measurements and analysis of malicious activities in critical infrastructures. In our future work, we endeavor to delve deeper in the data to find characteristics of (critical) sectors, such as device types and network services unique to a given sector, including the share of IoT or ICS/SCADA devices. Long-term monitoring of such presented and proposed characteristics would also be useful for analyzing threat trends threatening individual sectors of interest.

ACKNOWLEDGEMENT

The authors would like to express their sincere gratitude to the anonymous reviewers for their constructive feedback. This work was supported by a grant from the U.S. National Science Foundation (NSF) (Office of Advanced Cyberinfrastructure (OAC) #1755179) and by ERDF "CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence" (No. CZ.02.1.01/0.0/0.0/16_019/0000822).

REFERENCES

[1] Cristina Alcaraz and Sherali Zeadally. 2015. Critical infrastructure protection: Requirements and challenges for the 21st century. *International Journal of Critical Infrastructure Protection* 8 (2015), 53 – 66.

[2] Eray Balkanli and A. Nur Zincir-Heywood. 2014. On the analysis of backscatter traffic. In *39th Annual IEEE Conference on Local Computer Networks Workshops*. 671–678.

[3] Norbert Blenn, Vincent Ghiëtte, and Christian Doerr. 2017. Quantifying the Spectrum of Denial-of-Service Attacks Through Internet Backscatter. In *Proceedings of the 12th International Conference on Availability, Reliability and Security (ARES '17)*. ACM, New York, NY, USA, Article 21, 10 pages.

[4] Roland Bodenheim, Jonathan Butts, Stephen Dunlap, and Barry Mullins. 2014. Evaluation of the ability of the Shodan search engine to identify Internet-facing industrial control devices. *International Journal of Critical Infrastructure Protection* 7, 2 (2014), 114 – 123.

[5] Elias Bou-Harb, Mourad Debbabi, and Chadi Assi. 2014. Cyber Scanning: A Comprehensive Survey. *IEEE Communications Surveys Tutorials* 16, 3 (Third 2014), 1496–1519.

[6] Center for Applied Internet Data Analysis. 2018. UCSD Network Telescope – Near-Real-Time Network Telescope Dataset. http://www.caida.org/data/passive/telescope-near-real-time_dataset.xml. (2018). Accessed 2018-03-05.

[7] Thomas M. Chen and Saeed Abu-Nimeh. 2011. Lessons from Stuxnet. *Computer* 44, 4 (April 2011), 91–93.

[8] Council of European Union. 2008. Council Directive 2008/114/EC of 8 December 2008 on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection. https://publications.europa.eu/en/publication-detail/-/publication/ba51b03f-66f4-4807-bf7d-c66244414b10/. *Oficial Journal of the European Union* (8 December 2008). Accessed 2018-03-05.

[9] Ang Cui and Salvatore J. Stolfo. 2010. A Quantitative Analysis of the Insecurity of Embedded Network Devices: Results of a Wide-area Scan. In *Proceedings of the 26th Annual Computer Security Applications Conference (ACSAC '10)*. ACM, New York, NY, USA, 97–106.

[10] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. ACM, New York, NY, USA, 542–553.

[11] Jesse M. Ehrenfeld. 2017. WannaCry, Cybersecurity and Health Information Technology: A Time to Act. *Journal of Medical Systems* 41, 7 (24 May 2017), 104.

[12] European Commission. 2018. Critical Infrastructures. https://ec.europa.eu/home-affairs/what-we-do/policies/crisis-and-terrorism/critical-infrastructure_en. (14 March 2018). Accessed 2018-03-05.

[13] European Union Agency for Network and Information Security. 2015. Methodologies for the identification of Critical Information Infrastructure assets and services. https://www.enisa.europa.eu/publications/methodologies-for-the-identification-of-ciis. (23 February 2015). Accessed 2018-03-05.

[14] European Union Agency for Network and Information Security. 2016. Stocktaking, Analysis and Recommendations on the protection of CIIs. https://www.enisa.europa.eu/publications/stocktaking-analysis-and-recommendations-on-the-protection-of-ciis. (21 January 2016). Accessed 2018-03-05.

[15] Claude Fachkha, Elias Bou-Harb, Anastasis Keliris, Nasir Memon, and Mustaque Ahamad. 2017. Internet-scale Probing of CPS: Inference, Characterization and Orchestration Analysis. To appear.

[16] Claude Fachkha and Mourad Debbabi. 2016. Darknet as a Source of Cyber Intelligence: Survey, Taxonomy, and Characterization. *IEEE Communications Surveys Tutorials* 18, 2 (Secondquarter 2016), 1197–1227.

[17] Mario Galluscio, Nataliia Neshenko, Elias Bou-Hard, Yongliang Huang, Nasir Ghani, Jorge Crichigno, and Georges Kaddoum. 2017. A First Empirical Look on Internet-scale Exploitations of IoT Devices. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*.

[18] Juan David Guarnizo, Amit Tambe, Suman Sankar Bhunia, Martin Ochoa, Nils Ole Tippenhauer, Asaf Shabtai, and Yuval Elovici. 2017. SIPHON: Towards Scalable High-Interaction Physical Honeypots. In *Proceedings of the 3rd ACM Workshop on Cyber-Physical System Security (CPSS '17)*. ACM, New York, NY, USA, 57–68.

[19] Alexander Khalimonenko, Oleg Kupreev, and Kirill Ilganaev. 2018. DDoS attacks in Q4 2017. https://securelist.com/ddos-attacks-in-q4-2017/83729/. (6 February 2018). Accessed 2018-03-05.

[20] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions, and Defenses - Volume 9404 (RAID 2015)*. Springer-Verlag New York, Inc., New York, NY, USA, 615–636.

[21] Robert Lipovsky. 2017. Seven years after Stuxnet: Industrial systems security once again in the spotlight. https://www.welivesecurity.com/2017/06/16/seven-years-stuxnet-industrial-systems-security-spotlight/. (16 June 2017). Accessed 2018-03-05.

[22] John Materly. 2009. Shodan. https://shodan.io. (2009). Accessed 2018-03-05.

[23] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. 2015. IoTPOT: Analysing the Rise of IoT Compromises. In *9th USENIX Workshop on Offensive Technologies (WOOT 15)*. USENIX Association, Washington, D.C.

[24] Lukas Rist, Johnny Vestergaard, Daniel Haslinger, and Adrea De Pasquale. 2018. Conpot ICS/SCADA Honeypot. http://conpot.org/. (2018). Accessed 2018-03-05.

[25] The U.S. Department of Homeland Security. 2017. Critical Infrastructure Sectors. https://www.dhs.gov/critical-infrastructure-sectors. (2017). Accessed 2018-03-05.

[26] Emmanouil Vasilomanolakis, Shreyas Srinivasa, Carlos Garcia Cordero, and Max Mühlhäuser. 2016. Multi-stage attack detection and signature generation with ICS honeypots. In *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*. 1227–1232.

[27] Emmanouil Vasilomanolakis, Shreyas Srinivasa, and Max Mühlhäuser. 2015. Did you really hack a nuclear power plant? An industrial control mobile honeypot. In *2015 IEEE Conference on Communications and Network Security (CNS)*. 729–730.

[28] Eric Wustrow, Manish Karir, Michael Bailey, Farnam Jahanian, and Geoff Huston. 2010. Internet Background Radiation Revisited. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC '10)*. ACM, New York, NY, USA, 62–74.

[29] Eileen Yu. 2016. Majority CEOs unwilling to share cybersecurity information with outsiders. http://www.zdnet.com/article/majority-ceos-unwilling-to-share-cybersecurity-information-with-outsiders/. (17 February 2016). Accessed 2018-03-05.