

Mind Your Own Business: A Longitudinal Study of Threats and Vulnerabilities in Enterprises

Platon Kotzias

IMDEA Software Institute &
Universidad Politécnica de Madrid
platon.kotzias@imdea.org

Leyla Bilge and Pierre-Antoine Vervier

Symantec Research Labs
{leyla_bilge, pierre-antoine_vervier}@symantec.com

Juan Caballero

IMDEA Software Institute
juan.caballero@imdea.org

Abstract—Enterprises own a significant fraction of the hosts connected to the Internet and possess valuable assets, such as financial data and intellectual property, which may be targeted by attackers. They suffer attacks that exploit unpatched hosts and install malware, resulting in breaches that may cost millions in damages. Despite the scale of this phenomenon, the threat and vulnerability landscape of enterprises remains under-studied. The security posture of enterprises remains unclear, and it's unknown whether enterprises are indeed more secure than consumer hosts. To address these questions, we perform the largest and longest enterprise security study so far. Our data covers nearly 3 years and is collected from 28K enterprises, belonging to 67 industries, which own 82M client hosts and 73M public-facing servers.

Our measurements comprise of two parts: an analysis of the threat landscape and an analysis of the enterprise vulnerability patching behavior. The threat landscape analysis studies the encounter rate of malware and PUP in enterprise client hosts. It measures, among others, that 91%–97% of the enterprises, 13%–41% of the hosts, encountered at least one malware or PUP file over the length of our study; that enterprises encounter malware much more often than PUP; and that some industries like banks and consumer finances achieve significantly lower malware and PUP encounter rates than the most-affected industries. The vulnerability analysis examines the patching of 12 client-side and 112 server-side applications in enterprise client hosts and servers. It measures, among others, that it takes over 6 months on average to patch 90% of the population across all vulnerabilities in the 12 client-side applications; that enterprise hosts are faster to patch vulnerabilities compared to consumer hosts; and that the patching of server applications is much worse than the patching of client-side applications.

I. INTRODUCTION

Despite all efforts of the cyber security community, malware and other cyber attacks run rampant on the Internet. In recent years, there is almost not a single day we do not come across new incidents, such as data breaches [5] and ransomware attacks [7], [8]. Such incidents typically involve malware and affect both enterprises and consumers. While the security posture of consumers against malware and other cyber threats has been explored by security vendors [19], [29], [43] and the academic community [21], [32], the security posture

of enterprises against those threats has been significantly under-studied. This is problematic because enterprises own a significant fraction of the hosts connected to the Internet and possess valuable assets, such as financial data and intellectual property that may be the objective of (targeted) attacks. Enterprises may differ from consumers in important ways, such as using the same software across hosts, establishing security policies, installing multiple security products, educating their employees, and having departments dedicated to securing their assets. However, there exists a large variety of enterprises in terms of size, industries they belong to, financial assets, and security investment. Thus, it is very likely that the best practices mentioned above do not equally apply to all of them.

Currently, it is not clear how the security posture of enterprises differs according to different factors and whether enterprises are indeed more secure than consumers, i.e., if their security investment is paying off. In this paper, we aim to shed light into these questions by conducting what is, to the best of our knowledge, the largest and longest measurement study of enterprise security. Our data covers nearly 3 years and is collected from 28K enterprises with over 82M client hosts and 73M public-facing servers. We analyze the enterprise threat landscape including the prevalence of malware and PUP in enterprise client hosts and how common security practices, such as vulnerability patching and operating system updates are handled. We use a wealth of datasets collected from a large cyber security company (Symantec) and public sources. At the core of our study are file reputation logs that capture the installation of files in 82M real enterprise client hosts. These logs enable us an *internal view* of enterprise security. We complement these logs with a classification of the 28K enterprises into 67 industries and with AV labels of low-reputation files for classification. To analyze the security of the externally visible enterprise servers, we supplement our internal view, with an *outside-in view* using data from public sources, such as Internet-wide scans and blacklists.

Most related to our work is a study by Yen et al. [48] on the security of a large multinational enterprise comprising of 85,000 hosts for which they had four months of logs from an AV engine deployed in the enterprise. Similar to that work, we have an internal view of enterprise security, but our study analyzes a time frame that is eight times longer and covers 28K enterprises across 67 industries. Other related works have studied the network hygiene and security posture of enterprises using an outside-in view based on Internet-wide scans and blacklists [16], [27], [28], [50]. The limitations of an outside-

in view is that it only applies to externally reachable servers or is based on coarse-grained blacklists. Thus, its ability to provide an accurate view of the enterprise security posture remains to be proven. In contrast, we only use the outside-in view to complement and compare with our internal view of the enterprises. We find that the enterprise threat landscape looks very different from the inside and from the outside.

This paper comprises two main parts: an analysis of the enterprise threat landscape and an analysis of the enterprise vulnerability patching behavior. The first part of the paper, on analyzing the enterprise threat landscape, studies the encounter rate of malware and PUP in enterprises. It examines low reputation files installed on enterprise client hosts; classifies them into malware and PUP families; measures their prevalence; identifies the top families overall and per industry; examines whether some families target specific industries; analyzes the temporal evolution of the encounter rate; performs a case study on ransomware; and finally analyzes the outside-in view of enterprises by cross-checking blacklists with the publicly-facing IP addresses (including externally-facing servers).

The second part of the paper consists in the analysis of the vulnerability patching behavior of enterprises. We measure the time needed to patch 50% and 90% of the vulnerable population for 12 client-side applications installed on the 82M enterprise client hosts and 112 services installed on the 73M enterprise servers. For this, we first identify the list of vulnerabilities and vulnerable versions for those applications using NVD [33]. Then, we examine the time when those vulnerable versions are updated using the file reputation logs for client applications and Internet-wide scans for server applications. We also rank industries based on their patch deployment agility. Prior work has performed a similar study on client applications installed on consumer hosts [32] and has analyzed specific server vulnerabilities (e.g., Heartbleed) and misconfigurations [14], [26], [45]. However, to our knowledge, we are the first to measure the patch deployment behaviour of such a large number of enterprises, and to combine both client-side and server-side perspectives.

Below we list the most significant findings of our study:

- Between 91% (conservative estimate) and 97% (lax estimate) of the enterprises, 13% and 41% of the client hosts respectively, encountered at least one malware or PUP over the length of our study. Thus, despite their differences almost all enterprises will encounter some malware or PUP in three years.
- The 10 most-affected industries have 69%–76% of hosts affected, while the 10 least-affected have 15%–36%, highlighting that some industries, e.g., banks and finance-related, are definitely doing better than others.
- 73% of the low reputation files installed on enterprise hosts are unknown to VirusTotal, despite many being high prevalence. This questions how representative VirusTotal data may be of the enterprise landscape.
- Enterprises encounter malware (34% lax) much more often than PUP (8% lax). This is in contrast to prior works on consumer hosts that have shown that 54% had some PUP installed [21].

Dataset	Data	Count
File Reputation Logs 04/2015 – 12/2017	Hosts	82.1 M
	Enterprises	28 K
	Countries	137
	Total Reports	375 B
	Total Distinct Files	326 M
	Low Reputation Files	14.6 M
File Appearance Logs	Hosts	23M
	Enterprises	25 K
VirusTotal	Reports	1.3 M
NVD 01/2015 – 12/2017	Client Apps	12
	Client CVE	1,850
	Server Apps	112
	Server CVE	988
Internet Scans 10/2015 – 11/2017	Protocols	8
Blacklists 07/2015 – 12/2017	IP and Domain blacklists	38
Enterprise-to-IP mapping 07/2015 – 12/2017	Enterprises	28 K

TABLE I: Summary of datasets used.

- Cracking tools for Microsoft products (e.g., KM-SPico [20]) are found on 34% of all enterprises.
- Despite its notoriety, we observe ransomware affecting only a modest 0.02% of all enterprise client hosts.
- It takes over 6 months on average to patch 90% of the population across all vulnerabilities in the 12 client-side applications. This shows that patching still remains an issue even in enterprise settings.
- Enterprise computers are faster to patch vulnerabilities compared to consumer hosts.
- The patching of servers is overall much worse than the patching of client applications. On average a server application remains vulnerable for 7.5 months. Furthermore, it takes more than nine months for 90% of the enterprise server population to be patched.

II. DATASETS

This section details the datasets used in our work, summarized in Table I. We use *file reputation logs* to identify malicious files installed in 82M hosts across 28K enterprises; *file appearance logs* to identify the installation of 12 benign applications in the enterprises; *enterprise classification* to place enterprises into industries; *VirusTotal* (VT) reports to obtain AV labels to classify the malicious files; the *National Vulnerability Database* (NVD) to identify vulnerabilities in client-side and server-side applications and the range of versions affected; *Censys* [1] Internet-wide IPv4 scans to analyze externally-facing servers in the enterprises; *blacklists* to identify compromised hosts in the enterprises; and *enterprise-to-IP mapping* to check ownership of IP addresses.

File reputation logs. These logs capture metadata about the presence of files in 82M Windows client hosts across 28K enterprises in 137 countries and their corresponding reputation scores. These logs are collected from real client hosts in use

by enterprise customers of the cyber security company. The enterprises opted-in to sharing their data and the hosts and enterprises are anonymized to preserve the privacy of the customers. The dataset covers nearly three years from April 2015 to December 2017.

Each host in the collection regularly queries a centralized system to obtain the reputation of files installed in the host. The query includes file metadata such as file hash, file size, and publisher (if the file is signed). The response includes a reputation score that ranges between 128 and -127 with higher (positive) scores indicating good reputation and lower (negative) scores indicating lack of trust. The reputation score is computed using input from different security products and covers a large variety of features including file characteristics, dynamic behaviours, file prevalence, download source, and signer information.

We use this dataset to analyze the presence of malicious files in the enterprises. To identify malicious files, we first select a subset of 14.6M low reputation files (out of the 326M reported files) with a reputation score less than -20 . We selected this threshold experimentally to minimize the number of benign files included, while balancing the amount of data to be processed. The low reputation files can be of different types including executables (.exe, .dll, .sys), documents (.pdf, .docx), and archives (.zip, .rar). Overall, out of the 375B reports in the dataset, the low reputation files appear in 135M reports.

In detail, each report in the file reputation logs contains a timestamp, an anonymized enterprise identifier, an anonymized host identifier, a SHA256 file hash, and the file path where the file was installed. For each file hash, the logs also contain reputation score, AV detection label (if the file was flagged as malicious by the cyber security company), and file signer subject and code signature validation result (if the file is signed). For each anonymized host identifier, the logs also contain the Windows version installed (i.e, major and minor OS and Service pack versions).

File appearance logs. These logs capture metadata about all executables and archives installed in 23M real hosts belonging to 25K enterprises, a subset of the hosts and enterprises in the file reputation logs. Each event in the dataset can correspond to a download of an executable file or a compressed archive over the network, or the extraction of an executable file from a compressed archive. File appearance logs are very similar to file reputation logs, but differ in that they include all files installed in the host (regardless of their reputation score or potential maliciousness), they are collected from a smaller set of security products, they only include executables and archives, and they provide a more accurate timestamp of the first appearance of the file in the host. Again, the enterprises opted-in to sharing their data and the hosts and enterprises are anonymized to preserve privacy. The file appearance logs contain a timestamp, anonymized enterprise and host identifiers, SHA2 file hash, file signer, file path, and file version fields. We use these logs to identify the presence of specific versions of 12 selected benign client applications in the enterprise hosts.

Enterprise classification. Each anonymized enterprise identifier has associated its industry, number of employees, and country they are based in. This information was obtained from

Industry	Ent.	Hosts	IPs	Emp.	CC
Banks	1.1K	16.6M	7.6M	5.5M	85
IT Services	1.0K	7.5M	3,500M	3.0M	52
Healthcare Providers	1.1K	6.5M	2.9M	2.3M	46
Professional Services	875	3.8M	374.1M	1.4M	39
Commercial Services	1.2K	3.2M	366.5M	2.0M	49
Insurance	597	3.2M	1.4M	1.2M	52
Capital Markets	851	2.0M	4.1M	596K	55
Software	832	2.0M	803.8M	497K	43
Electronic Equipment	1.0K	1.7M	304.1M	1.7M	45
Machinery	1.4K	1.5M	13.3M	1.6M	49
Specialty Retail	601	1.5M	17.9M	1.6M	51
Construction & Engineering	1.3K	1.1M	471.9K	1.3M	52
Media	971	1.5M	96.6M	1.3M	44
Chemicals	850	1.0M	1.3M	909K	54
Food Products	846	872K	594.0K	1.6M	61
Financial Services	602	827K	749.1K	317K	47
Hotels Restaurants & Leisure	567	752K	1.2M	2.8M	46
Trading Companies	718	714K	12.4M	542K	40
Internet Software & Services	567	572K	407.8M	207K	34
Metals & Mining	874	506K	1.9M	1.8M	56

TABLE II: Number of enterprises, hosts, IPv4 addresses, employees, and country codes for the top 20 industries sorted by number of hosts. The high number of IPs for IT Services is due to that industry including ISPs and hosting providers.

an specialized external company. The classification comprises of 67 industries. Table II shows the number of enterprises, hosts, IP addresses, employees, and country codes for the top 20 industries by number of hosts in the file reputation logs. These top 20 industries cover 65% of the hosts. Banking is the top industry with 16.6M hosts across 1.1K banks in 85 countries, followed by IT services, and healthcare providers. Overall, the dataset shows good industry coverage with 55 (82%) of the industries having at least 100 enterprises and over 100K hosts.

VirusTotal. We query the hash of low reputation files in VirusTotal [47] (VT), an online service that analyzes files and URLs submitted by users using a large number of security tools. VT offers a commercial API that given a file hash returns metadata on the file including the list of detection labels assigned by a large number of AV engines used to scan the file. We use the AV labels as input to our malicious file classification. Unfortunately, given the API restrictions, we are only able to collect VT reports for 1.3M low reputation executables, corresponding to 18% of the 7.3M executables found among the 14.6 low reputation files.

NVD. We use the National Vulnerability Database (NVD) [33] to obtain the list of vulnerabilities, found between April 2015 and December 2017, in the selected benign client and server applications, For each vulnerability, we use the NVD to obtain the list of application versions affected by the vulnerability.

Internet scans. To identify vulnerabilities on servers belonging to the enterprises, we use data from IPv4 Internet-wide scans from Censys.io [1]. The scans were performed on multiple ports between October 2015 and November 2017. We use raw protocol banners from FTP, SSH, SMTP, IMAP(S), POP(S), and HTTP(S) scans. We extract application names and versions from these banners and match them against NVD data to identify vulnerable servers.

IP and domain blacklists. We also identify compromised hosts inside the enterprises using archives of 38 public and commercial IP and domain reputation blacklists. These blacklists include, among others, Abuse.ch [9], Cymru’s botnet tracking feeds [4], DShield [18], Phishtank [3], ShadowServer [46], Spamhaus DNSBLs [44], and Uceprotect DNSBLs [6]. These blacklists capture different types of malicious behaviors from clients and servers including, among others, spam, botnet infections, malicious server hosting, and brute-force login attacks. Each blacklist is downloaded on a hourly or daily basis depending on its update policy. The archives span 2.5 years between July 2015 and December 2017.

Enterprise-to-IP mapping. To analyze blacklisted hosts and vulnerabilities in externally-facing enterprise servers, we need to identify the public IP addresses an enterprise uses. These include IP addresses allocated to the enterprise, as well as IP addresses leased from cloud hosting providers. We obtain the blocks of IP addresses allocated to an enterprise from the Internet Routing Registries (IRRs). To identify the cloud hosting infrastructure used by an enterprise, we first use domain Whois data to identify domains that have been registered by the company or its subsidiaries. Then we use Rapid7’s passive DNS [37] to identify IP addresses that those domains have resolved to. Any IP address that is also a target for a domain from a different enterprise is removed to prevent pollution. To minimize the impact of IP address churn, we recompute the whole enterprise-to-IP mapping every week using archives of the data sources that cover the analysis period. We match blacklists and network scans with the enterprise-to-IP mapping for the corresponding week. We have verified the correctness of our mapping by manually validating it for 100 companies. We selected companies of different sizes and industries to account for potentially different IT administration practices.

A. Selection Bias

Our datasets may introduce selection bias. First, they only include enterprises investing in security products. Enterprises with no security products should have a worse security posture, making our results conservative. Also, our datasets only cover enterprises with security products of a specific vendor and that opted-in to share their data. Products from other vendors may provide different security, and enterprises that opted-out due to privacy concerns could be more security conscious. Furthermore, the file reputation and appearance logs contain only Windows client hosts. Client hosts running other OSes (e.g., macOS, Android) may have a different security posture. To analyze enterprise servers we use blacklists and network scans, but may miss internal servers not facing the Internet.

B. Ethical and Privacy Considerations

The file reputation logs and file appearance logs were collected from enterprises who opted in to sharing their data. Those logs are anonymized to preserve the privacy of the enterprises and their users. They do not contain any identifiable data about the origin of the log entries. Machines and enterprises are referred to using anonymized machine and enterprise identifiers. The outside-in analysis requires the list of enterprise customers of the cyber security company to identify their external-facing IP addresses. That analysis was

Data	Count
Low Reputation Files	14.6M
Low Reputation Executables	7.3M
Benign Files	729K
Total VT reports collected	1.3M
Executables with vendor label	3.3M
Total Labeled	2.0M
Total families	19K
Malware Families	15.5K
PUP Families	3.5K

TABLE III: Breakdown of low reputation files.

performed by an employee of the cyber security company and the customer list was not shared with the external authors. To further prevent deanonymization of the enterprises and their users, we present our findings on an aggregated level and on anonymized case studies.

III. THREAT LANDSCAPE

This section presents our analysis of the enterprise threat landscape. We start by analyzing the security posture of enterprise client hosts from inside (Sections III-A through III-E). First, we describe our family classification of malicious files in Section III-A. Then, we analyze the prevalence of malware and PUP (Section III-B) and how specific families are to industries and enterprises (Section III-C). Next, we perform a longitudinal analysis of malware and PUP encounters (Section III-D) and present a case study on the prevalence of ransomware in the enterprises (Section III-E). Finally, we analyze the security posture of enterprises from the outside (including externally-facing servers) in Section III-F.

A. Family Classification

To analyze the most prevalent threats enterprise client hosts face, we identify the malicious files in our dataset and classify them into families. We start with 14.6M low-reputation files described in Section II. We first filter out the benign files that might have been assigned a low reputation, e.g., due to their low prevalence. This step removes executables signed by benign publishers – using a manually curated whitelist of 948 popular publishers – as well as executables for which a VT report is available and are considered malicious by less than 4 AV engines. As a result, we filter out 729K executables.

Table III summarizes our classification. Out of the 7.3M executables among the low-reputation files, we collected 1.3M VT reports. Not all 7.3M files were queried to VT due to API restrictions. Among those queried, VirusTotal only knew 27%. This is important because our community largely assumes VT data adequately represents the malware ecosystem.

Our threat classification methodology analyzes the AV labels in the 1.3M VT reports, as well as the labels assigned by the cyber security company, available for another 3.3M executables. We feed the AV labels as input to AVClass [39]. AVClass outputs the most likely family name for each sample and also classifies it as malware or PUP based on the presence of PUP-related keywords in the AV labels (e.g., *adware*, *unwanted*). In addition to the files flagged as PUP by AVClass,

Family	Type	Hosts	Ent.	Files
opencandy	pup	1.1 M	19.5K	12.0K
winactivator	malware	470.8K	9.4K	5.3K
installcore	pup	453.4K	17.3K	54.6K
autoit	malware	398.4K	6.5K	12.2K
remoteadmin	pup	333.0K	8.7K	1.7K
sogou	pup	282.8K	2.2K	813
micrayment	pup	264.0K	3.1K	21
asparnet	pup	232.8K	13.7K	238
elex	pup	218.3K	7.1K	6.9K
donex	pup	179.0K	2.0K	49
dealply	pup	176.5K	12.8K	23.9K
nssm	malware	171.2K	441	41
ramnit	malware	142.8K	7.6K	737.2K
qjwmonkey	pup	142.3K	2.0K	281
asprox	malware	139.7K	2.1K	1.4K
flystudio	malware	126.9K	3.0K	5.7K
conficker	malware	125.6K	5.0K	2.4K
spigot	pup	114.2K	10.0K	1.3K
fusioncore	pup	111.2K	9.3K	901
ursu	malware	108.2K	2.3K	559

TABLE IV: Top 20 families by number of hosts.

we further identified PUP samples by matching their publisher information with 3.8K known PUP publishers. The original version of AVClass is designed to take as input VT reports, which include labels from multiple AV vendors. For the 3.3M executables for which we only have one AV label, we had to modify AVClass by removing the check that requires a family to appear in at least two AV engines to be considered. While no longer using a plurality vote for those 3.3M files, AVClass still enables us to remove noise and generic tokens from the cyber security company’s labels. Overall, we labeled 2.1M (29%) executables belonging to 19K families. For the remaining 2.5M samples for which labels were available, no family was identified as their labels were generic.

One advantage of our classification over prior works that classified malware obtained from malware feeds (e.g., [23]) is that we can rank malware families based on their prevalence on real hosts, while samples in malware feeds may be biased towards highly polymorphic families. Table IV shows the top 20 malware and PUP families the enterprises encountered over the analysis period. From the top 20 families, 12 are PUP and 8 are malware families. The most prevalent family is *opencandy*, a well-known commercial pay-per-install service [21], which we observe installed in 1.1M hosts in over 19K enterprises. The most popular malware family is *winactivator*, a label used by AVs for Microsoft Windows crack tools. Activators are found on 34% (9.4K) of all enterprises across all industries. These enterprises have a median size of 490 hosts, although there are also 98 large enterprises (over 100K hosts). Furthermore, 10% of these 9.4K enterprises had *winactivator* installed in over 15% of their client hosts. Further analysis reveals that the majority of the *winactivator* executables belong to KMSPico [20], a popular Microsoft Windows and Office crack tool. The publishers of KMSPico claim that the cracked software can get all the available updates by Microsoft.

Preval.		All	Mal.	PUP
Lax	Host	33.6M (41%)	28.2M (34%)	6.2M (8%)
	Ent.	27.2K (97%)	26.8K (96%)	24.8K (89%)
Con.	Host	10.8M (13%)	8.3M (10%)	5.2M (6%)
	Ent.	25.5K (91%)	24.5K (87%)	24.6K (87%)

TABLE V: Lax and conservative PUP and malware prevalence estimates.

B. Malware and PUP Prevalence

In this section, we analyze malware and PUP encounters in enterprises. We first establish their prevalence using a lax and a conservative estimate. The lax estimate measures the prevalence of all low reputation files minus the benign files, a total of 13.9M files. The conservative estimate measures the prevalence of only the executables for which we have a VT report and are not benign. Table V summarizes the prevalence results. We find that using the lax estimate 41% of the hosts and 97% of the enterprises have suffered at least one malware or PUP encounter during the nearly three years analyzed. Using the conservative estimate the numbers are 13% of hosts and 91% of the enterprises. Thus, regardless of the estimate used, the vast majority (91%–97%) of enterprises have suffered at least one malware or PUP encounter. Only 3%–9% of the enterprises never encountered malware or PUP in our analysis period. All these clean enterprises had less than 100 hosts and the vast majority had only one host. These findings highlight the difficulty of securing enterprises against malicious software. Any reasonable-sized enterprise can be expected to encounter malicious software in three years.

We compare our measured prevalence with prior works. Yen et al. [48] observed that 15% of hosts in a single enterprise encountered malware over a four-month period in 2014. Microsoft reports a 2017 malware encounter rate of 14% in Canada [43], however, without making the distinction between enterprise and consumer hosts. While our conservative estimate is close to those prevalence rates, our lax estimate shows a higher prevalence than those prior works.

The split between malware and PUP shows a higher impact of malware than PUP, with malware affecting 10%–34% of hosts and 87%–96% of enterprises, compared to 6%–8% and 87%–89% for PUP. These findings indicate that both malware and PUP affect the vast majority of enterprises, although malware impacts a larger number of hosts. We find that PUP encounters in enterprises are considerably lower of what is reported in previous works for consumer hosts. Kotzias et al. [21] measured PUP prevalence in consumer hosts for the period Jan 2013 - July 2014 and found out that 54% of the 3.9M analyzed hosts had had some PUP installed. There are four PUP families in Table IV whose prevalence was measured for consumer hosts in [21]. They all show much higher prevalence in consumer hosts: *opencandy* (8% in consumer hosts vs 1.3% in enterprise hosts), *installCore* (8.5% vs 0.55%), *dealply* (2.8% vs 0.2%), and *spigot* (2.6% vs 0.1%). This confirms that PUP is significantly less prevalent in enterprises. This could be due to stricter security policies about what programs can be installed applied by enterprises, which may affect PUP, but not malware (since PUP typically requires user acceptance for installation). Other explanations




Fig. 1: Number of families per host

Industry	All	Mal.	PUP
Electrical Equipment	76.4%	69.7%	22.0%
Automobiles	75.5%	70.4%	13.9%
Construction Materials	74.4%	66.7%	18.5%
Marine	74.3%	67.4%	30.8%
Semiconductors	72.9%	66.8%	19.9%
Industrial Conglomerates	72.8%	67.5%	26.0%
Communications Equipment	71.3%	63.8%	22.0%
Healthcare Equipment	70.8%	64.1%	15.6%
Leisure Products	69.4%	60.6%	11.5%
Beverages	69.3%	61.0%	10.9%
Thriffs and Mortgage Finance	36.5%	30.2%	8.3%
Diversified Financial Services	35.7%	30.3%	4.6%
Specialty Retail	35.3%	29.6%	6.4%
Healthcare Providers	34.0%	27.4%	3.1%
Professional Services	32.6%	27.2%	4.5%
Real Estate	31.8%	25.7%	2.4%
Wireless Telecommunication	28.6%	23.3%	6.6%
Biotechnology	20.5%	15.1%	1.1%
Consumer Finance	15.9%	11.4%	1.9%
Banks	15.7%	13.6%	1.2%

TABLE VI: Top 10 (most affected) and bottom 10 (least affected) industries by malware and PUP prevalence.

could be differences on the awareness of corporate users and the different time period of the two studies, as prior work shows PUP prevalence dropping at the end of 2015 [23].

Figure 1 shows the cumulative distribution of the number of distinct families observed per host with at least one encounter. Nearly 50% of the hosts are only affected by one family and 75% by less than 5 families. The fact that 25% of hosts encountered more than 5 families is surprising and can be due to pay-per-install relationships [13], [21] or to machines that are periodically re-infected.

Industry prevalence. Table VI presents the top ten (most-affected) and bottom ten (least-affected) industries ranked by malware and PUP prevalence. There is a significant difference between both groups of industries. The most-affected industries have 76%–69% of hosts affected, while the least-affected industries have 36%–15%. That is, the ten most-affected industries have more than twice the prevalence of malware and PUP compared to the least-affected ten industries.

This shows that there are industries that take security more seriously than others. Four of the ten least-affected industries are finance-related including Banks and Consumer Finance, which are the two least-affected industries. This matches reports that banking is the industry that invests the most in cyber security products [2]. However, note that Banks have the most hosts in our dataset, thus a prevalence of 16% represents over 2.5M encounters. The most-affected industries are Electrical Equipment, Automobiles, and Construction Materials. In general, this group seems dominated by industries related to manufacturing and consumer products.

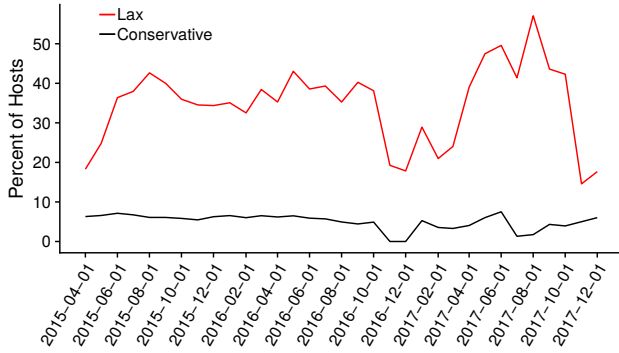
C. Malware and PUP Specificity Analysis

We rank the top 20 PUP and malware families affecting each industry, whose union comprises of 221 malware and 86 PUP families. We observe a few families that appear in the top 20 of almost all industries. These include both malware (*winactivator*, *ramnit*, *autoit*) as well as PUP (*dealply*, *installcore*, *spigot*, *amonetize*, *opencandy*, *asparnet*, *remoteadmin*). On the other hand, 117 of the malware families and 57 of the PUP families were not found on 90% of the industries. This is an interesting observation showing that there are many PUP and malware families only seen in one or a small number of industries. Furthermore, 17 malware families were found only on one industry and those families were on that industry’s top 3 malware list. For example, the remote access trojan *xtrat* is only in the top 20 malware of the Construction and Engineering industry, but encountered in 2% (22K) of those hosts.


We perform the same investigation at per-enterprise level to identify families targeting specific enterprises. The number of malware families seen in only one enterprise is 1,911 (37%), while for PUP is 446 (26%). Thus, the specificity of malware families is higher than for PUP families. On the other hand, in contrast with the industry-based results, we do not observe any malware or PUP family encountered in the majority of enterprises. Among the enterprise-specific malware families, 78 are ranked as the top malware family encountered in that enterprise, which may indicate targeting. Of those 78, 13 affect a large enterprise. One example is the *zcrypt* ransomware. It affected only one enterprise in our dataset (from the Hotels, Restaurants, and Leisure industry). In conclusion, we find a significant number of enterprise-specific malware families and observe indications of targeting for 78 families.

D. Longitudinal Analysis

Figure 2 shows the monthly encounter rate using the conservative and lax estimates. The percentage of hosts that encountered malware (Figure 2a) using the conservative estimate does not change drastically over the years, remaining on average around 7%. On the other hand, we observe larger fluctuations on the lax estimate. Between mid-2015 and the end of 2016, the monthly encounter rate was stable around 30–40%. Then, in November 2016 the percentage drops by approximately 20%. In May 2017, the monthly encounter rate increases drastically reaching over 50% in August. Finally, in November 2017 it drops by 25%. The two large drops one year apart are also visible in the percentage of enterprises encountering malware (Figure 2b). The reasons for these two large malware encounter drops remain unknown. We checked



(a) By hosts



(b) By enterprises.

Fig. 2: Monthly malware and PUP prevalence by number of hosts and enterprises with at least one encounter.

Family	Hosts	Ent.	Ind.	Files
wannacry	30.1K	872	65	2.2K
locky	20.3K	5.2K	67	4.6K
petya	11.2K	155	46	72
ransomkd	10.2K	1.1K	66	70
teslacrypt	9.4K	2.9K	66	5.9K
cryptolocker	8.7K	1.7K	66	714
cerber	6.1K	2.2K	66	1.7K
cryptowall	2.6K	1.4K	66	359
dcraptor	2.0K	468	59	36
torrentlocker	785	443	62	207
All	103K	8.8K	67	16K

TABLE VII: Top 10 ransomware families by number of hosts.

that the abrupt increase in mid-2017 is not due to *wannacry* and *petya* that emerged during that time (see Section III-E). However, it could be due to other malware families exploiting the same EternalBlue vulnerability.

E. Case Study: Ransomware

In this section we present a case study on the prevalence of ransomware in enterprise networks. There are 28 ransomware families among the classified low reputation executables. In total, we identify 103K hosts in 8.8K enterprises across all 67 industries affected by ransomware. This is a pretty low




Fig. 3: Monthly number of hosts and enterprises with ransomware appearances.

prevalence of 0.02%. Assuming an average ransom payment of \$500, these encounters amount to a modest \$51.5M in direct costs, plus possibly an order of magnitude larger indirect costs including remediation [10].

Table VII presents the top 10 ransomware families by number of affected hosts. *wannacry* and *locky* are the ransomware families found in most hosts, 30K and 20K, respectively. *wannacry* only ranks seventh in terms of enterprises affected due to the fact that its worm-like behavior exploited the previously known *Eternal Blue* SMB vulnerability that was patched for Windows 7 and above. Thus, it likely only affected enterprises with hosts using earlier Windows versions, but spread quickly within those enterprises. In fact, from the 12K enterprises with at least one Windows XP host, 50% experienced a ransomware attack, while the average encounter rate was only 31%.

Most ransomware families are found in the majority of industries, indicating that ransomware operators currently do not target specific industries. Figure 3 shows the monthly number of hosts and enterprises with ransomware encounters in our analysis period. We observe the first large peak on March 2016 where affected hosts reach 7K and affected enterprises 3K. This peak is mostly due to *locky*. The second and largest peak occurs on June 2017 and is due to *wannacry* and *petya*. It affects more than 13K hosts, but (as explained above) the number of affected enterprises does not significantly increase. Note however that this peak in June 2017 is much smaller than the one observed in Figure 2 for all malware encounters. Thus, these two families by themselves cannot explain that large increase.

F. Outside-in Perspective

In this section, we look at the enterprise threat landscape from an *outside-in* perspective. We extract symptoms of malware encounters inside an enterprise by (i) uncovering all public, Internet-facing IP addresses owned or used by an enterprise and (ii) correlating them with datasets of external indicators of compromise (IoCs), essentially blacklists of Internet hosts associated with different types of malicious activity (e.g., spam senders, C&C servers of known botnets, malware distributing and phishing web servers).

Industry	Ent.	Botnet Infect.	Brute-f. Logins	C&C Host.	Malware Distri.	Misc. Attacks	Phish.	Scan.	Spam	Total
Media	598	108.8K	1.8K	3.5K	30.0K	34.7K	6.9K	949	610.7K	797.2K
Communications Equipment	141	13.6K	55	763	2.8K	11.5K	1.3K	14	456.5K	486.5K
Software	543	7.5K	1.8K	3.4K	9.2K	26.3K	6.2K	581	336.5K	391.5K
Technology Hardware	129	5.0K	1.4K	2.3K	8.4K	30.2K	5.2K	533	222.4K	275.4K
IT Services	675	6.8K	1.8K	8.1K	25.0K	71.0K	13.0K	194	94.9K	220.7K
Internet Software and Services	371	7.0K	1.9K	3.1K	16.8K	18.1K	13.7K	361	85.8K	146.9K
Electronic Equipment	563	2.0K	46	1.8K	4.7K	7.6K	3.4K	13	116.8K	136.3K
Diversified Consumer Services	118	8.5K	75	574	4.6K	1.7K	849	43	90.9K	107.2K
Commercial Services	756	502	156	2.6K	4.5K	9.4K	3.8K	16	35.0K	56.0K
Professional Services	614	764	159	3.4K	7.4K	10.9K	5.0K	14	10.5K	38.2K
Multi-Utilities	17	194	3	51	101	55	59	0	688	1.2K
Airlines	47	20	2	119	268	185	226	0	329	1.1K
Construction Materials	100	6	0	170	263	187	260	0	263	1.1K
Paper and Forest Products	92	14	1	132	255	169	215	0	252	1.0K
Transportation Infrastructure	80	14	2	53	156	84	125	0	194	628
Multiline Retail	13	58	0	23	53	34	53	1	127	349
Gas Utilities	34	2	0	46	85	55	68	0	76	332
Marine	38	4	0	33	70	46	60	0	86	299
Water Utilities	25	1	0	40	68	46	61	0	68	284
Tobacco	5	2	0	45	55	40	53	0	62	257

TABLE VIII: Breakdown of malicious activity exhibited by industries (top 10 and bottom 10).

Table VIII presents the breakdown of the malicious activity observed from the top and bottom 10 industries in terms of number of blacklisted IP addresses. The first big trend we can observe from the blacklisted hosts inside enterprises is that, as of today, *spam* is still the predominant type of malicious activity sourced by allegedly compromised machines. In most companies of most industries, spam largely dominates any other type of malicious activity. This phenomenon can be in part explained by the fact that spam is heavily monitored and might be easier to detect than machines hosting malware or C&C servers inside an enterprise. The high prevalence of *malware*, *phishing* and *C&C server hosting* highlights the serious threat that compromised machines inside enterprises can pose.

Comparing the top and bottom 10 industries from the perspective of malware encounters (Table VI) and blacklisted hosts (Table VIII), we notice some obvious differences. Only one industry – Communications Equipment – appears both in the top 10 malware industries and the top 10 blacklisted hosts industries. Moreover, two industries – Construction Materials and Marine – are found in the top 10 malware industries and in the bottom 10 blacklisted hosts industries. There are two likely reasons for this: (i) blacklists have limited visibility into much malware encountered in enterprise client hosts, and (ii) most malware do not exhibit external IoCs captured by blacklists.

The effectiveness of blacklists for operational threat detection has already been extensively studied in previous works [22], [23], [31], [35], [41], [42]. Some of these studies have also assessed the quality of blacklists. The general take-away message from these prior studies is that *the quality of blacklists can vary drastically from one another* so care should be taken when selecting them, *more specialized datasets should be preferred, when available*, and that *despite their limitations, they remain a useful source of malicious activity*. As we have seen in our results here-above, IP- and domain-based blacklists can be useful to provide a general trend on the security posture

of an enterprise and, by extension, the industry it belongs to. However, to study malware encounters in enterprises, we can see that blacklists cannot match the granularity and accuracy of more specialized datasets, such as the file appearance logs used in Section III. We understand that, in the absence of other datasets, blacklists may be the only source to study the (enterprise) threat landscape. However, care should be taken when deriving conclusions solely based on blacklists.

IV. VULNERABILITY PATCHING BEHAVIOR

In this section we analyze the presence of vulnerabilities and their patching behavior in enterprise networks, which prior work has shown to be fairly correlated with future security incidents [12], [27]. In particular, we study vulnerability patching practices carried out by different industries. Our goal is to understand the security posture of the enterprises and whether particular industries are less or more agile to patch their vulnerabilities and therefore, are less or more secure against cyber threats. We conclude the section with an analysis of OS upgrade behavior in enterprises.

We analyze both vulnerabilities in client and server applications. In Section IV-A we use the file appearance logs to analyze the patching speed of 12 popular client-side applications in the enterprise hosts. In Section IV-B we use periodic IPv4 Internet-wide scans to analyze the patching speed of vulnerabilities in 112 server applications and libraries. Since these are externally-facing servers installed in the enterprises, they are easier to be discovered by the attackers and are greatly exposed to external threats.

A. Analysis of client-side vulnerabilities.

Our analysis of client-side vulnerabilities focuses on 12 client applications and frameworks: .NET, Adobe Air, Adobe Reader, Chrome, Firefox, Internet Explorer, Java Runtime Environment (JRE), MariaDB, Silverlight, Skype, Thunderbird,

TABLE IX: Client application patching summary. It shows the number of application versions, the number of hosts and enterprises where the application was installed, the number of vulnerabilities analyzed, the number of hosts unpatched at the end of the analysis, the 50% and 90% enterprise patch time in days measured in this work, and the 50% and 90% consumer patch time in days measured in previous work [32].

Program	Vendor	Versions	Hosts	Enterprises	CVE	Unpatched	Enterprise PT		Consumer PT [32]	
						Hosts	50%	90%	50%	90%
Chrome	Google	267	10.2M	23,814	454	1.7M	18	78	15	246
Firefox	Mozilla	205	4.2M	20,575	308	1.1M	25	161	36	179
Thunderbird		10	159K	6,132	40	15K	23	98	27	129
Skype	Microsoft	41	1.1M	18,120	2	8K	17	89	-	-
Internet Explorer		1,035	15.8M	24,543	428	11M	47	138	-	-
.NET		197	8.5 M	22,474	21	2.5M	60	162	-	-
Silverlight		43	8.9M	22,763	17	5M	82	182	-	-
Media Player		141	9.5M	22,700	1	7.5M	147	314	-	-
JRE	Oracle	340	5.7M	22,367	21	1.4M	56	141	-	-
Air	Adobe	11	1.2M	15,136	316	216K	44	152	-	-
Reader		47	13.9M	23,563	221	6.2M	78	234	188	219
MariaDB	-	35	13.5K	1,106	53	3K	75	246	-	-
	TOTAL	2,372	23M	25,367	1,882	AVG	67	200		

and Windows Media Player. We selected these 12 applications because they are popular; they cover both stand-alone applications (e.g., Chrome, Adobe Reader) and frameworks (.NET, JRE); they include proprietary programs from five vendors (Adobe, Google, Microsoft, Mozilla, Oracle) and one open-source application (MariaDB); their executables are signed; and they embed the program version in their executables.

To identify the presence of these applications on the enterprise hosts, we follow a methodology similar to that proposed by Nappa et al. [32]. We first identify the main executable for each application (e.g., firefox.exe), then we examine the file appearance logs to obtain the hashes and file versions of all executables with that name and signed by the right publisher (e.g., Mozilla). This step outputs for each application, a mapping from file hash to the application version corresponding to that hash. Using this mapping we can identify hosts in the file appearance logs where those versions were installed, as well as their installation time. We then use the NVD to obtain the vulnerabilities, disclosed between April 2015 and December 2017, in those 12 applications and the list of vulnerable program versions for each vulnerability.

For each vulnerability we compute the *patch time*, i.e., the time needed to patch a certain fraction (50% and 90% in this work) of the vulnerable hosts. To compute the patch time, we exclude hosts that never patched a vulnerability, e.g., because they left the population.

Table IX summarizes the client application patching results. The left side of the table captures, for each application, the name, the vendor, the number of versions identified, the number of hosts with one of those versions installed, the number of enterprises those hosts belong to, and the number of vulnerabilities analyzed. Overall, we analyze 1,882 vulnerabilities, of which 50% are critical (CVSS ≥ 9) and 90% have high impact (CVSS ≥ 7). The most popular application is Internet Explorer installed in 69% of the hosts in the file appearance logs, followed by Adobe Reader (60%), and Chrome (44%). Eight of the 12 applications are installed in over 20K enterprises highlighting their popularity.

The middle part of the table summarizes our enterprise

patching measurements. It shows the average number of hosts that never patched and the average time in days to patch 50% and 90% of the vulnerable hosts. The results show that Chrome is the fastest application being patched requiring on average 18 days to patch 50% of the vulnerable hosts and 78 days to patch 90%. On the other hand, the slowest application is Windows Media Player which takes nearly 5 months to patch 50% of the vulnerable hosts and over 10 months to patch 90%. Overall, it takes over 6 months on average to patch 90% of the population across all applications and vulnerabilities, highlighting the limitations of patch deployment in enterprises.

The right side of the table shows the patch time reported by Nappa et al. [32] in their analysis of 8.4M consumer hosts. We use a similar methodology to that work and examine four applications in common: Chrome, Firefox, Thunderbird, and Adobe Reader. The comparison shows that three of the four applications (Chrome, Firefox, Thunderbird) reach 90% patching faster in enterprises and another three (Firefox, Thunderbird, Reader) reach 50% patching also faster in enterprises. These results seem to indicate that enterprises are on average faster to apply patches than consumers. One caveat is that the period of analysis differs between both works, 2008–2013 for the work on consumer hosts and 2015–2017 in this work. There may be different reasons behind the improvement in patching in enterprise hosts including enterprises being more security aware, having deployed security software that may detect the need to update, having teams dedicated to securing their hosts, or that enterprise hosts may be online more often than consumer hosts (enabling the patches to be downloaded earlier).

Client patching by industry. Table X ranks the top and bottom ten industries by vulnerability patching time. We provide detailed patching time (50% and 90%) for the five client applications that are installed the most on enterprise hosts. We also provide results averaged across all applications. The results are obtained by cumulating the ranking for each application and ordering the list over the cumulative ranking. As it can be seen, the industries that invest the most in cyber security products, and encounter higher amount of malware

TABLE X: Industry ranking of vulnerability patching time (in days).

Rank	Industry	IE		Chrome		Adobe Reader		Firefox		JRE		All Apps	
		50%	90%	50%	90%	50%	90%	50%	90%	50%	90%	50%	90%
1	Communications Equipment	42	91	20	71	85	201	20	111	73	148	53	152
2	Consumer Finance	45	123	18	67	72	193	20	126	67	156	52	152
3	Diversified Financial Services	46	142	14	67	71	191	27	119	50	140	56	164
4	Diversified Telecommunication Services	48	134	17	70	90	247	22	130	30	127	49	141
5	Capital Markets	47	120	22	48	81	228	25	133	77	157	64	159
6	Software	47	140	15	66	63	196	20	118	51	155	55	160
7	Trading Companies and Distributors	50	138	17	76	88	210	23	155	63	104	56	152
8	IT Services	42	114	18	78	69	214	24	152	31	109	60	156
9	Health Care Technology	46	147	18	78	62	172	23	144	61	149	49	133
10	Diversified Consumer Services	46	124	16	67	64	185	29	159	78	195	65	151
57	Containers and Packaging	54	143	14	74	93	300	27	188	53	167	58	183
58	Multi-Utilities	54	136	17	74	71	470	58	306	63	159	62	210
59	Road and Rail	47	132	16	86	95	266	32	204	42	161	64	184
61	Real Estate Management and Development	50	177	18	79	80	295	22	163	39	162	55	178
62	Textiles, Apparel and Luxury Goods	52	187	16	76	78	259	24	177	58	183	63	178
63	Industrial Conglomerates	58	201	15	75	82	281	22	172	78	177	65	196
64	Air Freight and Logistics	52	174	20	80	95	371	31	240	63	150	58	185
65	Gas Utilities	60	187	22	93	108	256	31	179	70	162	68	197
66	Construction Materials	49	169	18	100	107	341	36	189	84	193	66	187
67	Multiline Retail	60	276	15	78	55	256	32	251	88	219	61	193

in count (not in percentage), such as finance, software and communications are considerably faster at patching their vulnerable applications. On the other hand, the industries whose majority of machines encounter malware are worst at patching their vulnerabilities on a timely manner making the window of their exposure to cyber threats larger. Seeing industries such as gas and electricity utilities in the bottom part of the list is especially worrisome as successful attacks in this kind of industries could have physical impacts. When we perform a similar analysis on the percentage of unpatched hosts and the length of their vulnerability windows in each industry, we obtain different rankings. While 90% of the machines from the top best (i.e. Banks, Household Products, Multi-Utilities) industries remain vulnerable for an average of four months, hosts from the bottom of the list (i.e. Tobacco, Multiline Rail, Energy Equipment and Services, and Marine) remain vulnerable for 15 months.

Disabling automatic updates. We examine whether enterprises may have disabled the auto-update functionality of applications, which is mandatory but can be disabled through configuration options. For this, we compute the average time it takes each host to install a new version of applications (for all versions, not only the vulnerable ones). Then we examine the distribution across hosts to identify outlier hosts that update applications slower. For this, we first calculate the median of each application’s update speed distribution. To identify outliers, we calculate the absolute deviation which was proposed as an optimal way for outlier detection [25]. Using the absolute deviation, hosts that on average take more than $median + absolute_deviation$ days to update their apps are considered outliers. We then look for enterprises for which the majority of hosts are outliers, which would indicate an enterprise-wide policy to disable auto-updates. We only find a limited number of enterprises that satisfy that condition. For example, for Chrome we found two, for Adobe Reader four and for Firefox only one enterprise where more than 75% of the machines were outliers. Thus, disabling auto-updates on the client applications we analyzed is in general a rare policy.

Best and worst patchers We identify the best and the worst patchers in our data to compare their malware encounter rate with their patching behavior. We choose enterprises that have at least 1000 machines for this measurement. The top 10 enterprises that patch their vulnerable applications the fastest patch 90% of their machines in less than 10 days. Note that here we take the average patch time for all of the applications we analyzed in our study. On the other side of the scale, the 90% patch time of the worst patchers is 500 days on average. While the best patcher is an enterprise from the Hotels, Restaurants and Leisure industry that patches most hosts in only 5 days, the remaining best patchers are from the Financial and Insurance industry. The worst ones are from the Capital Market, Media, Speciality Retail, Textiles, Apparel and Luxury Goods and Healthcare. Having the worst patcher from an industry which was ranked as the 5th best in patching and the best patcher from an industry that is ranked as the 4th worst illustrates the big variation in patching behavior among companies. We also looked at the malware prevalence in these enterprises and found out that the worst patchers encounter more malware compared to the best patchers. This simple investigation on the best and worst patchers supports that patching applications on time has a significant effect on the number of malware encounters.

B. Analysis of server-side vulnerabilities

In this section we analyze the patching of vulnerabilities in servers belonging to the 28K enterprises. Each server corresponds to an IP address and may run multiple services on different ports. Each service is an instance of one of the 112 server software packages analyzed. To identify the specific software and version of a service, we use a set of 2,664 regular expressions that are applied on the protocol banners collected through Internet-wide scans. One difference with the client application analysis is that here we do not know the exact timestamp when a service was updated. Instead, we approximate it with the time when the new version is first observed, which happens later as scans take place with at most daily granularity.

In the remainder of this section we use the same metrics to measure the patching behavior of enterprise servers than those we used for client machines, i.e., the *patch time* measuring the time it takes for a vulnerable server application to be updated once a patch is released, and the *vulnerability window* defining the time period during which a server application remains vulnerable to a known vulnerability.

Property	Count	Avg	50%	90%
Servers	73.1M	-	-	-
Vul. servers	17.9M	-	-	-
Patched at least once	16.4M	-	-	-
Never patched	1.5M	-	-	-
CVEs/server	-	6.82	5.00	15.00
CVE CVSS score	-	5.42	5.00	7.60

TABLE XI: Summary of the server-side applications vulnerability assessment. These results are computed for the 28 K enterprises and the 112 server-side applications.

Overview. Table XI provides an overview of the server vulnerability analysis results. These results are computed for the 28K enterprises and the 112 server-side applications. Out of 73.1M servers mapped to the 28K enterprises, 17.9M have had at least one vulnerable service. On average, each server is affected by more than six vulnerabilities. Even more worryingly, at least 10% of vulnerable servers are affected by more than 15 vulnerabilities. One important observation is that 1.5M servers in 11,905 enterprises have never been upgraded throughout the 2.5 years analysis period.

Server patching by application. Table XII presents the top 10 vulnerable server software in terms of number of vulnerable servers found. The table is dominated by popular SSH and Web servers. The top 3 server software (OpenSSH, Apache, IIS) had at least one vulnerable version installed on over 2.5M servers across more than 10K enterprises. On average, for the 112 server programs, it takes eight weeks (56 days) to patch 50% of the servers and over nine months (282 days) to patch 90%. While the server patch time for 50% is slightly shorter (56 days) that the patch time for 50% of the client applications (67 days), when considering 90% of the servers, it is almost 50% worse than the 200 days (90%) observed on the 12 client applications. One possible reason for the slower server patching is the lack of automatic updates on server programs. There is a stark contrast between the 50% and 90% patch time. While seven of the top 10 software have a 50% patch time of 24 days or lower (significantly better than the average), their 90% patch time is 10 months (the average). Thus, even for the most popular server software it is hard for enterprises to fully deploy patches. Finally, only 2.2K out of 28K (7.9%) enterprises have a 50% patch time below or equal to the average 50% patch time across all applications (56 days). While a significant fraction of the enterprises are diligent in patching their servers, the rest are quite slow making it very hard to completely eliminate a vulnerability. This situation creates points of entry for cyber-criminals to penetrate corporate networks by leaving Internet-facing servers vulnerable for very long periods of time.

Server patching by industry. We now focus on the patching behavior of enterprises per industry. Table XIII presents the




Fig. 4: Percentage of monthly enterprise hosts per Windows OS version.

top and bottom 10 industries based on their overall server-side service patch time. Overall, the patching behavior of enterprises in the top 10 industries, i.e., the best patchers, is not good, especially when compared to the per-client application patch time reported in Table IX. For instance, we can see that the 50% patch time across all applications for the top 10 industries (between 78 and 199 days) is way above the average 50% patch time across all applications (56 days). The same applies to the 90% patch time per industry (between 412 and 709 days) when compared to the average 90% patch time across all applications (282 days). Looking at the patch time per protocol, we can see that some very popular services like Web servers and SSH servers bear the worst patch time (both 50% and 90%) in the top and bottom 10 industries. Indeed, the top 10 industries take almost six months (174 days) to patch 50% of their SSH servers and almost seven months (205 days) to patch 50% of their web servers. This is more than three times the average 50% patch time (56 days) reported across all applications. Similar to what we observe for client-side applications, we witness some industries associated with critical infrastructures, such as Gas Utilities, Transportation Infrastructure, and Marine among the worst-patching industries. Overall, our conclusion is that the patching behavior of servers in enterprises is worryingly bad, across all server applications and services.

C. Operating System Upgrade Behavior

In this section, we analyze the Windows upgrade behavior in the enterprise client hosts. Figure 4 shows the monthly percentage of hosts that use Windows XP, Vista, 7, 8, and 10. There is no big change on Windows version usage between 2015 and 2017. For most of the period Windows 7 dominated with over 80% of the hosts using them. On March 2017, an increase of Windows 10 hosts occurs, raising from 10% to 20% by end of 2017. A simultaneous drop of Windows 7 machines indicates a slow shift from Windows 7 to 10. The percentage of Windows Vista and 8 remains constantly below 10%; there is no significant adoption of these versions by enterprises. Windows XP usage is already low in the beginning of 2015 (around 10%) and declines until the end of 2017.

Microsoft ended support of Windows XP in April 2014,

TABLE XII: Summary of the server-side applications and patching behavior of the enterprise servers. Results per application are given for the top 10 vulnerable applications in number of affected servers. The average, 50%, 90% patch time and the average vulnerability window are also provided for the total 112 server-side applications.

Rank	Program	Service	Vulnerable			Patch Time			Avg. Vul. Window
			Machines	Ent.	CVEs	Avg.	50%	90%	
1	OpenSSH	SSH	4,517,497	10,764	84	96	22	317	132
2	Apache Httpd	HTTP	2,691,805	10,655	182	108	24	323	165
3	Microsoft IIS	HTTP	2,690,361	13,738	22	140	32	552	208
4	Lighttpd	HTTP	1,133,379	908	26	78	15	233	88
5	vsftpd	FTP	825,480	2,045	5	59	7	216	89
6	mini_httpd	HTTP	810,859	349	2	89	15	253	111
7	Nginx	HTTP	413,911	4,890	14	175	162	346	191
8	ProFTPD	FTP	266,929	2,427	27	70	7	287	106
9	Apache Coyote	HTTP	208,213	2,834	1	168	71	575	241
10	Exim	SMTP	52,260	1,867	13	135	16	480	211
Total 112 Apps.						108	56	282	230

TABLE XIII: Industry ranking of server-side applications vulnerability patching time (in days). Blank fields indicate industries in which a server-side application was not found.

Rank	Industry	Machines	FTP		SSH		SMTP		HTTP(S)		POP(S)		IMAP(S)		All Apps	
			50%	90%	50%	90%	50%	90%	50%	90%	50%	90%	50%	90%	50%	90%
1	Multi-Utilities	1.0K	18	98	183	435			70	322	340	340			78	412
2	Communications Equipment	77.8K	57	433	155	543	127	705	155	695	377	532	319	535	159	679
3	Thriffs and Mortgage Finance	262	119	274	211	492	69	236	200	698					162	705
4	Beverages	596	130	619	316	695	188	724	169	668					171	695
5	Automobiles	4.0K	64	473	237	561	209	510	188	678	546	695	179	280	172	659
6	Technology Hardware	889.9K	137	621	155	540	226	390	183	660	169	629	39	593	172	660
7	Electric Utilities	22.8K	109	657	218	590	124	392	200	582	182	468	114	205	181	589
8	Multiline Retail	222	74	481	114	342	127	127	190	714	557	557			190	709
9	Food and Staples Retailing	2.6K	32	623	295	603	176	436	184	705	176	310			196	705
10	Internet Software and Services	665.4K	155	674	174	574	134	595	200	674	196	716	148	588	199	674
58	Construction Materials	196	151	496	310	543	188	226	234	590	134	303	141	141	226	610
59	Electrical Equipment	1.5K	134	736	285	579	58	472	200	599	328	479	550	550	230	606
60	Internet and Catalog Retail	2.8K	137	428	157	543	71	520	238	705					233	704
61	Containers and Packaging	1.1K	36	369	317	691	188	473	181	614	408	670			237	691
62	Gas Utilities	114	323	469	277	543	226	226	200	589					252	582
63	Construction and Engineering	2.6K	49	417	317	691	210	399	200	614	169	348			253	685
64	Personal Products	356	90	287	284	683	226	480	260	660			185	185	268	671
65	Energy Equipment and Services	264	22	399	317	513	241	350	203	630	78	306	74	74	279	625
66	Transportation Infrastructure	399	82	357	317	695	151	304	272	689	297	554	372	667	279	703
67	Marine	156	120	175	317	487	155	452	299	686	203	203	264	478	292	691

but we still see an alarmingly large number of enterprises that use it. During 2017, we see a total of 466K Windows XP hosts in more than 43% (12.2K) of the enterprises. Most of these are medium to large enterprises; 73% of those have a total of more than 100 hosts, and 25% more than 1K hosts. All 67 industries have at least some companies with outdated OS hosts. The three industries with the largest number of XP hosts are *Electronic Equipment, Instruments and Components, Specialty Retail*, and *Banks*. Interestingly, banks have the lowest percentage of hosts with malware appearances but still more than 500 of those operate Windows XP hosts. This possibly indicates the difficulty of decommissioning legacy systems.

We see far less Windows Vista hosts, compared to XP hosts, in enterprises during 2017; a total of 86K hosts in 7K enterprises. Microsoft ended support for Vista in April 2017. The low number of hosts is probably due to the small adoption of Windows Vista by enterprises (Figure 4). As in the case of XP, enterprises that still use Vista are medium to large. In fact, 76% (5.5K) of enterprises with Vista hosts have also XP hosts.

V. RELATED WORK

Threat landscape reports and studies To our knowledge, there are not many scientific works that performed systematic investigations on the cyber threat landscape of enterprises. This was mainly due to the absence of data that is representative and accurate enough for malware encountered by enterprises. The only way researchers could estimate the maliciousness of enterprises was to use public blacklists that provide information about known infected IP addresses [16], [50]. While the enterprise landscape is greatly understudied by the scientific community, there are many industrial annual threat reports [19], [29], [43]. These threat reports mainly focus on general statistics about malware seen on the Internet without making the distinction of the industry and consumer data. After targeted attacks towards specific industries hit the news in 2010, these reports started to provide industry-based statistics, however only on companies that encountered spear-phishing attacks [19].

Depending on which threats were popular on the particular year, these reports provide special details and create new sections that did not exist in previous years. For example,

in 2018, we see extensive details about ransomware due to wannacry and petya events in 2017. While the content slightly changes, some of the sections are consistent over the years. It is typically to list the top malware families for each year, the top zero-day and normal vulnerabilities. In our work, we also provide similar statistics however focusing on the threat landscape of enterprises and their vulnerability patching behavior which is shown to be significantly correlated with future malware infections [12]. By conducting this study, our goal was to understand whether particular enterprise profiles have weaker security practices and therefore, more attention should be paid to them to fix these issues before they become the next target. One important finding we found in the course of this study was that the industries that operate with critical infrastructures are very slow to patch their applications, making them vulnerable against possible future cyber attacks.

Securing enterprises Another line of research that relates to our work conducts studies to identify enterprise specific threat detection techniques [24], [34], [49] and protection mechanisms that rely on hardening the networks [30]. Levin et al. proposed to deploy Honeynets inside enterprise networks that typically have higher bandwidth usage and it is harder to detect malicious traffic [24]. The core idea here was that Honeynets are not suppose to send or receive any traffic, and anything that is observed in these are good indicators for malicious traffic and could be used to identify other infected machines. Yen et al. on the other hand aimed at improving the incident detection rate by mining security logs that are produced by various security products in an enterprise [49]. Similarly, Oprea et al. mined large-scale log data to identify enterprise infections at earlier stage. The key insight of the work was that the detection of early-stage infections could be modeled with belief propagation algorithms. The data is used for the experimentation was anonymized DNS logs and web proxy logs collected from a large enterprise. McDaniel et al. approaches the problem of securing enterprises differently by proposing to apply hardening policies [30]. The ambitious goal of the paper was to define the normal behavior such that anything else could be blocked. In this direction, the authors present techniques to automatically generate host profiles based on their historical interactions. In 2016, Edwards et al. released a report [16] where they found positive correlation between the presence of some “risky” externally-facing services, e.g., peer-to-peer file sharing, and misconfigured servers, e.g., vulnerable HTTPS servers, in enterprises and botnet infections inferred from blacklists.

Predicting the future. Recently, a number of works proposed various malware prediction methodologies that could be applied to the enterprise scenario. The closest scientific work to ours in the prediction domain [48] conducted a study on malware encounters in a single enterprise with the goal of discovering features that are highly correlated with future incidents. The investigation carried out by the paper includes the analysis of malware infection vectors, the network configuration details of the hosts when the likelihood of malware finding them is higher, and the correlation analysis of the user characteristics on receiving more malware such as the location of the user and the job details. Bilge et al. proposed a prediction method that can make predictions about future incidents as well [12]. The difference of this work from the

previous works is that it analyzed internal security telemetry of 18 enterprises to identify a long list of features that are good predictors for future malware infections. While this work does not provide statistics about the data that is available for us to compare, we got motivated from the fact that vulnerability patching behavior is one of the good predictors for good and bad security hygiene and decided to preform an investigation on the topic.

On the prediction topic, there are other works that aimed at finding correlating features that could be used for predictive analytics using publicly available datasets [27], [28], [50]. Zhang et al. measured the correlation between symptoms of network mismanagement, such as, the presence of open DNS resolvers, BGP misconfigurations, the presence of open mail relays, and externally observed malicious activity, such as, spam bots, botnet infections [50]. By combining results from Internet-wide scans with IP- and domain-based reputation feeds, such as, spam blacklists, they claim an overall strong positive correlation between network mismanagement and maliciousness. Furthermore, authors mention that gap that exists between such a statistical correlation and a validated causality. Liu et al. attempted to predict security incidents, such as, data breaches, website defacements, by correlating externally observed indicators of malicious network activity with security incidents reported by organizations [28]. The assumption behind their technique is that malicious activity originating from an organization’s network, measured using IP- and domain-based reputation feeds, e.g., spam blacklists, is somehow indicative of the security diligence of the organization. They build a machine learning-based model over about one year of data and achieve 60% TPR and 20% FPR. Finally, the same authors extended their previous work [27], still relying upon the assumption that the diligence of organizations to properly secure their network is correlated with externally observed indicators, they use a machine learning-based model to correlate (i) symptoms of network mismanagement observed at organizations around the Internet, e.g., untrusted HTTPS certificates, open DNS resolvers, etc, (ii) traces of malicious activity originating from these organizations, e.g., spam bots, botnet infections, etc, and (iii) security incidents publicly reported by these organizations. They can achieve up to 90% TPR with 10% FPR. Furthermore, they observe that network mismanagement-related features appear to be the best indicator of future security incidents. Getting motivated by all these exciting studies and their findings, we explored some of these features in a more detailed manner to understand how the security landscape of enterprises differs in various industries and the best security practices taken by them.

Vulnerability studies. Throughout this paper, we also present results about existing server and client-side vulnerabilities in different industries and the vulnerability patching behavior. Therefore, we will also present a brief related work about some recent works on the topic. Despite being one of the most studied topics, to our knowledge, there is no study that measured vulnerability patching behavior of client-side applications on enterprise computers. There is an extensive list of works that particularly focused on the vulnerability life-cycle topic [11], [17], [40] and analyzed the manual-patch deployment [36], [38].

Thanks to Durumeric et. al. who devised a technique to

perform Internet-wide scans in less than an hour [15], it became possible to know which server-side applications and their corresponding versions are deployed on enterprises and measure their vulnerability patching behavior. In a follow-up study, the same authors [14] measured the patching speed of the Heartbleed vulnerability in OpenSSL, finding out that more than half of the servers were still vulnerable after three months. Leveraging the Internet-wide scans from Zmap, Li et al. studied misconfigured anonymous FTP servers in [45] and discovered that no less than 20K of these servers leak sensitive data and provide cyber-criminals with easy takeover targets. Subsequently, in [26] Li et al. took advantage of large-scale scans to study the effectiveness of reporting misconfigured industrial control systems (ICS), IPv6 firewalls and DDoS amplifiers.

Finally, Nappa et al. investigated the patch deployment behavior of 8M users analyzing 1.5K vulnerabilities on 10 client-side applications [32]. One of the interesting findings of the paper was that to patch 90% of the population it took more than a year for most of the applications.

VI. CONCLUSION

In this work we have performed the largest and longest enterprise security study up to date using telemetry from 28K enterprises for nearly 3 years. In the first part of the work, we have analyzed the enterprise threat landscape finding that 91%–97% of the enterprises, and 13%–41% of the enterprise hosts, encountered at least one malware or PUP file over the length of our study; that 73% of low reputation files installed on enterprise hosts are unknown to VT; that enterprises encounter malware much more often than PUP; and that some industries like banks and consumer finances are doing notoriously better, achieving a three times lower malware and PUP encounter rates than the most-affected industries. We also assess the effectiveness of blacklists to infer the network hygiene of enterprises and conclude that such data sources fail to provide the granularity and magnitude of the various threats enterprises are exposed to.

In the second part of the work, we have analyzed patch deployment in enterprises. We measure that it takes over 6 months on average to patch 90% of the population across all vulnerabilities in 12 client-side applications. This shows that patching still remains an issue even in enterprise settings. However, we observe that overall enterprises are faster at patching their vulnerable hosts than consumer users. Finally, we observe that the patching of servers is overall much worse than the patching of client applications. It takes more than nine months for 90% of the enterprise server population to be patched. Both the client and server patching show that the vulnerability window is large enough for cybercriminals to exploit them and find their way into the corporate networks.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments and feedback. The research leading to these results has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreements No. 786669 and No. 731535. This work also was supported by the US National Science Foundation under

grants CNS-1528156, CNS-1564329, and OAC-1348077; by the UK EPSRC under grants EP/M013472/1, EP/K035584/1, and EP/P009301/1; by the Regional Government of Madrid through the N-GREENS Software-CM project S2013/ICE-2731; and by the Spanish Government through the DEDETIS grant TIN2015-7013-R. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors or originators, and do not necessarily reflect the views of the sponsors.

REFERENCES

- [1] Censys.io. <https://censys.io/>.
- [2] Investment in cyber security by industry sector. <https://www.savoystewart.co.uk/blog/firms-investment-on-cyber-security-by-industry>.
- [3] Phishtank. <https://www.phishtank.com/>.
- [4] Team Cymru. <http://www.team-cymru.com/>.
- [5] Top 15 data breaches in 2018. <http://www.businessinsider.fr/us/data-breaches-2018-4>.
- [6] Uceprotect DNSBLs. <http://www.uceprotect.net/>.
- [7] Petya ransomware outbreak: Here’s what you need to know. <https://www.symantec.com/blogs/threat-intelligence/petya-ransomware-wiper>, 2017.
- [8] Ransom.wannacry. <https://www.symantec.com/security-center/writeup/2017-051310-3522-99>, 2017.
- [9] Abuse.ch. Fighting malware and botnets. <http://www.team-cymru.com/>.
- [10] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. J. G. van Eeten, M. Levi, T. Moore, and S. Savage. Measuring the Cost of Cybercrime. In *Workshop on the Economics of Information Security*, San Francisco, CA, USA, May 2013.
- [11] W. A. Arbaugh, W. L. Fithen, and J. McHugh. Windows of vulnerability: A case study analysis. *IEEE Computer*, 33(12), December 2000.
- [12] L. Bilge, Y. Han, and M. Dell’Amico. Riskteller: Predicting the risk of cyber incidents. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS ’17, 2017.
- [13] J. Caballero, C. Grier, C. Kreibich, and V. Paxson. Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Proceedings of the 20th USENIX Security Symposium*, San Francisco, CA, USA, August 2011.
- [14] Z. Durumeric, J. Kasten, D. Adrian, J. A. Halderman, M. Bailey, F. Li, N. Weaver, J. Amann, J. Beekman, M. Payer, and V. Paxson. The matter of Heartbleed. In *Proceedings of the Internet Measurement Conference*, Vancouver, Canada, Nov 2014.
- [15] Z. Durumeric, E. Wustrow, and J. A. Halderman. Zmap: Fast internet-wide scanning and its security applications. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, pages 605–620, Washington, D.C., 2013. USENIX.
- [16] B. Edwards, J. Jacobs, and S. Forrest. Risky Business: Assessing Security with External Measurements. <https://bjedwards.github.io/assets/papers/EdwardsRisky.pdf>, August 2016.
- [17] S. Frei. *Security Econometrics: The Dynamics of (In)Security*. PhD thesis, ETH Zürich, 2009.
- [18] Internet Storm Center. DShield DNSBL. <https://www.dshield.org/>.
- [19] Internet Security Threat Report, 2018. <https://www.symantec.com/security-center>.
- [20] KMSPico. <http://kmspico10.com/>.
- [21] P. Kotzias, L. Bilge, and J. Caballero. Measuring PUP Prevalence and PUP Distribution through Pay-Per-Install Services. In *USENIX Security Symposium*, August 2016.
- [22] M. Kührer, C. Rossow, and T. Holz. Paint it Black: Evaluating the Effectiveness of Malware Blacklists. In *Proceedings of the 17th International Symposium on Research in Attacks, Intrusions and Defenses*, September 2014.

- [23] C. Lever, P. Kotzias, D. Balzarotti, J. Caballero, and M. Antonakakis. A Lustrum of Malware Network Communication: Evolution and Insights. In *Proceedings of the 38th IEEE Symposium on Security and Privacy*, San Jose, CA, USA, May 2017.
- [24] J. Levine, R. LaBella, H. Owen, D. Contis, and B. Culver. The Use of Honeynets to Detect Exploited Systems across Large Enterprise Networks. In *IEEE Information Assurance Workshop*, 2003.
- [25] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4):764 – 766, 2013.
- [26] F. Li, Z. Durumeric, J. Cxyz, M. Karami, M. Bailey, D. McCoy, S. Savage, and V. Paxson. You’ve got vulnerability: Exploring effective vulnerability notifications. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 1033–1050, Austin, TX, 2016. USENIX Association.
- [27] Y. Liu, A. Sarabi, J. Zhang, P. Naghizadeh, M. Karir, M. Bailey, and M. Liu. Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 1009–1024, Washington, D.C., 2015. USENIX Association.
- [28] Y. Liu, J. Zhang, A. Sarabi, M. Liu, M. Karir, and M. Bailey. Predicting Cyber Security Incidents Using Feature-Based Characterization of Network-Level Malicious Activities. In *Proceedings of the 2015 ACM International Workshop on International Workshop on Security and Privacy Analytics, IWSPA ’15*, pages 3–9, New York, NY, USA, 2015. ACM.
- [29] Threats Report by McAfee Labs, 2018. <https://www.mcafee.com/enterprise/en-us/assets/infographics/infographic-threats-report-mar-2018.pdf>.
- [30] P. D. McDaniel, S. Sen, O. Spatscheck, J. E. van der Merwe, W. Aiello, and C. R. Kalmanek. Enterprise Security: A Community of Interest Based Approach. In *Network and Distributed Security Symposium*, 2006.
- [31] L. Metcalf and J. M. Spring. Blacklist Ecosystem Analysis: Spanning Jan 2012 to Jun 2014. In *Proceedings of the 2Nd ACM Workshop on Information Sharing and Collaborative Security, WISCS ’15*, pages 13–22. ACM, 2015.
- [32] A. Nappa, R. Johnson, L. Bilge, J. Caballero, and T. Dumitras. The attack of the clones: A study of the impact of shared code on vulnerability patching. In *IEEE Symposium on Security and Privacy*, pages 692–708, 2015.
- [33] National vulnerability database. <https://nvd.nist.gov/>.
- [34] A. Oprea, Z. Li, T.-F. Yen, S. H. Chin, and S. Alrwais. Detection of Early-Stage Enterprise Infection by Mining Large-Scale Log Data. In *Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2015.
- [35] A. Pitsillidis, C. Kanich, G. M. Voelker, K. Levchenko, and S. Savage. Taster’s Choice: A Comparative Analysis of Spam Feeds. In *Proceedings of the 2012 Internet Measurement Conference, IMC ’12*, pages 427–440. ACM, 2012.
- [36] T. Ramos. The laws of vulnerabilities. In *RSA Conference*, 2006.
- [37] Rapid7 Labs. Forward DNS (FDNS). https://opendata.rapid7.com/sonar.fdns_v2/.
- [38] E. Rescorla. Security holes... who cares. In *Proceedings of the 12th USENIX Security Symposium*, pages 75–90, 2003.
- [39] M. Sebastian, R. Rivera, P. Kotzias, and J. Caballero. Avclass: A tool for massive malware labeling. In *Research in Attacks, Intrusions, and Defenses*, pages 230–253, 2016.
- [40] M. Shahzad, M. Z. Shafiq, and A. X. Liu. A large scale exploratory analysis of software vulnerability life cycles. In *Proceedings of the 2012 International Conference on Software Engineering*, 2012.
- [41] C. A. Shue, A. J. Kalafut, and M. Gupta. Abnormally Malicious Autonomous Systems and Their Internet Connectivity. *IEEE/ACM Trans. Netw.*, 20(1):220–230, Feb. 2012.
- [42] S. Sinha, M. Bailey, and F. Jahanian. Shades of Grey: On the Effectiveness of Reputation-based “blacklists”. In *Proceedings of the 3rd International Conference on Malicious and Unwanted Software (MALWARE ’08)*, pages 57–64, Fairfax, Virginia, USA, October 2008.
- [43] Microsoft Security Intelligence Report, 2018. https://info.microsoft.com/rs/157-GQE-382/images/EN-US_CNTNT-eBook-SIR-volume-23_March2018.pdf.
- [44] Spamhaus. The Spamhaus Project. <https://www.spamhaus.org/>.
- [45] D. Springall, Z. Durumeric, and J. A. Halderman. FTP: The Forgotten Cloud. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 503–513, June 2016.
- [46] The ShadowServer Foundation. ShadowServer. <https://www.shadowserver.org/>.
- [47] VirusTotal. <http://www.virustotal.com/>.
- [48] T.-F. Yen, V. Heorhiadi, A. Oprea, M. K. Reiter, and A. Juels. An Epidemiological Study of Malware Encounters in a Large Enterprise. In *ACM SIGSAC Conference on Computer and Communications Security*, 2014.
- [49] T.-F. Yen, A. Oprea, K. Onarlioglu, T. Leetham, W. Robertson, A. Juels, and E. Kirda. Beehive: Large-Scale Log Analysis for Detecting Suspicious Activity in Enterprise Networks. In *Annual Computer Security Applications Conference*, 2013.
- [50] J. Zhang, Z. Durumeric, M. Bailey, M. Liu, and M. Karir. On the Mismanagement and Maliciousness of Networks. In *NDSS*, 2014.