

On Gaps in Enterprise Cyber Attack Reporting

Abulfaz Hajizada
School of Cyber Studies
The University of Tulsa
Tulsa, OK, USA
abh7867@utulsa.edu

Tyler Moore
School of Cyber Studies
The University of Tulsa
Tulsa, OK, USA
tyler-moore@utulsa.edu

Abstract—It has long been lamented that firms underreport cyber attacks. In recent years, regulators have begun mandating that certain organizations must publicly report when incidents occur. Adherence to these requirements is an empirical question that has been largely unexamined to date. In this paper, we study regulatory filings by U.S. public companies to the Securities Exchange Commission and to the Department Health and Human Services that discuss cyber attacks. We also compare the findings against crowdsourced reports of cyber incidents appearing in media outlets. We find substantial gaps in coverage, both in terms of attacks that make the news but do not appear in regulatory filings and vice versa. We conclude by discussing the implications for the study of cyber attack and defense as well as for policymakers.

1. Introduction

A longstanding challenge in cybersecurity research is gathering data on when attacks occur. Victims often do not like to broadcast to the world when they have been attacked, since it may harm their reputation and invite unwanted attention. Nonetheless, information on cyber attacks do often come to light for a variety of reasons. First, organizations may disclose an attack if its effects cannot be hidden, such as experiencing a ransomware attack that harms availability. Second, breach-notification laws (beginning with California’s law enacted in 2002) oblige organizations to disclose to consumers when private personal data is exposed. Third, regulators at the Securities and Exchange Commission have issued guidance calling on publicly-listed companies to disclose when cybersecurity breaches occur.

Researchers can in turn study cyber attack prevalence and evolution by gathering data from public sources. Indeed, a key motivation for regulations mandating disclosure is to reduce information asymmetries that could improve defender decision-making in how best to counter attacks [1], [2]. Unfortunately, issues with reporting complicate the collection and analysis of public data on cyber attacks. Each of the public sources has its own criteria for inclusion and methodology for collecting data. For example, the Privacy Rights Clearinghouse maintains a database of incidents that involve breaches of private information [3]. Similarly, the US Department of Health and Human Services publishes breaches of patient data by US health providers [4]. These services focus on data breaches

because laws compel disclosure when this category of attack occurs.

But what about the myriad other types of cyber attack? In 2011, publicly-traded companies received guidance from the US Securities and Exchange Commission encouraging them to disclose in their annual filings any cyber incidents of material significance to the company [5]. The guidance was clarified in 2018 [6] and culminated in a proposed rule in 2022 [7]. Consequently, publicly-traded firms are now expected to disclose all types of cyber attacks. The disclosures are written in free-form text in the filings submitted to the SEC. While promising greater coverage of additional attack types, the obligation is limited to publicly-traded firms and can be difficult to identify given the unstructured nature of the text in the filings.

Another source of cyber attack data comes from media reports. For example, The crowdsourced volunteer service Hackmageddon scours the Internet for reports of all types of cyber attacks and regularly publishes its findings [8]. Researchers at Temple University have focused on ransomware attacks [9], which also relies on public reporting. Since ransomware attacks affect service availability, many such attacks are disclosed by victims and naturally attract coverage.

Despite the plethora of sources, little is known about how these data sources relate to one another. For example, when a cyber attack is covered in the media, does the victim disclose its occurrence to the SEC? What about the other direction: do firms ever manage to bury attack reports in regulatory filings without being picked up in the media? Do certain types of attacks appear more often in regulatory filings or news reports than others? The goal of this paper is to dig deeper into these disparate sources to learn more about attack reporting and any gaps that exist across sources. Improving our understanding of the benefits and limitations of such reporting will help evaluate the effectiveness of disclosure as remedy to information asymmetries present in cybersecurity.

2. Related Work

Most empirical research studying cyber attacks on firms has focused on data breaches (e.g., [10]–[12]). Gay studied the relationship between news reports and breaches, similar to our comparison between media reports and regulatory filings [12]. He demonstrated that firms strategically time breach announcements to mitigate their impact. More recently, Li et

al. compared data breach reports to the 10-K regulatory filings of breached firms [13]. They utilized the regulatory filings to measure security awareness at firms while relying on data breach announcements from other sources.

Over the years, many cybercrime measurement researchers have studied gaps in attack datasets. Moore and Clayton studied phishing URL feeds, finding huge gaps in the private datasets maintained by website-takedown companies that explained why some phishing websites took far longer to remove than others [14]. Pitsillidis found similarly large gaps in email spam feeds [15], as have more recent investigations into threat intelligence feeds [16], [17]. These papers are similar in spirit to our work, which focuses on gaps in higher-level datasets of regulatory filings and news reports detailing cyber attacks. Nonetheless, we should not be surprised that the gaps persist in the context of the current study.

3. Data and Methodology

We now review the datasets utilized in the paper: SEC regulatory filings, patient data breaches reported to HHS, and media reports of cyber attacks collected by the volunteer group Hackmageddon. The final dataset can be viewed online at <https://www.dropbox.com/s/8qydpmwk8mwx74t/FinalCombinedSECResearchData.xlsx?dl=0>.

3.1. Dataset 1: SEC

Since 2011, firms listed at US stock exchanges have been encouraged to disclose significant cyber incidents in the discussion of risk factors on regulatory filings.

Using the SEC’s search engine EDGAR, we searched for all 10-K reports that included the term “cyber” anywhere within a document for the years 2017–2022. We had initially experimented with searching for specific terms, such as “ransomware” and “data breach”. However, we discovered through trial and error that some filings avoid using such terms. The more general “cyber” term was much more broadly used, perhaps taking the lead from the SEC’s own guidance that employs the term.

In total, we identified 7809 companies that mentioned cyber in at least one of their 10-K annual filings. The vast majority of these filings did not describe cyber attacks; rather, they discussed risk factors or other aspect of cyber unrelated to experiencing an attack.

For each 10-K, we automatically extracted text from the PDFs. We only considered “cyber” when mentioned inside the section on risk factors. We then identified the paragraph before and after the mention of the term “cyber” and then manually inspected the text to determine if it described an attack. When an attack was identified, it was further classified as describing a data breach, ransomware, or other attack.

We inspected a random sample of 1300 companies filing 10-Ks out of the total population of 7809 companies. From that sample, we identified 18 distinct attacks during the six year span of study.




Figure 1: Venn diagrams comparing coverage of attacks in SEC 10-K filings and Hackmageddon.

3.2. Dataset 2: HHS

Following the passage of the HITECH Act in 2010, health-care organizations have been required to disclose to the US Department of Health and Human Services when breaches of patient data occur. HHS in turn publishes data on its website. We gathered all reports of HHS data breaches between 2017 and 2022. In total, 170 health-care organizations are publicly-traded companies. Just 15 reports came from publicly-traded firms.

3.3. Dataset 3: Hackmageddon

Hackmageddon is a crowdsourced volunteer service that reports multiple types of cyber attacks from news reports. Its data is divided into single files that are sorted by individual months per year.

We examined the entire Hackmageddon dataset starting from 2017 through 2022, which averaged 1783 records per year. We then pared down this dataset significantly to include only relevant attacks and victims. First, we checked for the names of the companies to determine if they were publicly traded or not, keeping only those that were publicly traded. Second, we manually inspected the textual description of each attack in order to determine whether it targeted the company itself or individual customers of the company. We only included the former in our study. In other words, the incident had to specifically mention a company name or its corporate network. Attacks targeting customers, such as a news article reporting on Apple users falling for phishing attacks, are excluded.

In total, we identified 110 attacks matching our definitions affecting 103 distinct companies during the inspection period. Following this identification and labeling process, we inspected the SEC 10-K filings of all identified companies to determine whether they reported those particular attacks.

4. Results

In Section 4.1, we compare the coverage in identifying attacks among the three sources. In Section 4.2, we examine whether the type of attack reported varies by data source and over time. In Section 4.3, we examine the 10-K filings of attacked companies to determine whether the companies reported awareness to cyber risks before and after the incident occurs.

Health Data Breaches




Figure 2: Venn diagrams comparing coverage of attacks reported to HHS with SEC 10-K filings and Hackmageddon.

	Data Breach		Other Intrusion		Ransomware	
Hackmageddon	75	68%	9	8%	26	24%
SEC	12	60%	4	20%	4	20%
HHS	14	93%	0	0%	1	7%
Total	90	68%	13	10%	29	22%

TABLE 1: Attack types split by source

4.1. Comparing Data Sources

We first report on the sample of 1300 firms filings with the SEC. 20 of these reports disclosed a cyber attack. Of these, five also appeared in Hackmageddon. A further six firms from the sample were attacked according to media reports observed in Hackmageddon but did not mention these attacks in their annual filings with the SEC. We know this because we checked the SEC filings for all companies with attacks reported on Hackmageddon, regardless of whether their reports mentioned the term “cyber”. The results are illustrated in the left Venn diagram in Figure 1.

It is noteworthy that so many attacks are reported in only one source or the other. This indicates that there is significant underreporting by firms to the SEC, even when the attack has been made public by some other means. Moreover, most attacks reported to the SEC are not picked up in the media.

We can estimate the total number of attacks reported to the SEC between 2017–22. Assuming a similar proportion of attacks can be identified for the remaining 6509 firms whose filings have not yet been inspected, we estimate that 120 attacks were reported.

We next consider the combined dataset of the SEC sample and all Hackmageddon reports. In total, 110 distinct attacks on publicly-traded companies are identified by Hackmageddon. Less than half the time (46), the affected firms mention the cyber attack in their 10-K filings. Once again, this confirms that there is substantial underreporting of attacks to the SEC. We estimate that a further 90 attacks were reported to the SEC but not identified in public reports by Hackmageddon.

We next examine the attacks involving patient health information reported to HHS. Because most healthcare in the US is not provided by publicly-traded firms, the dataset is relatively

Attack Types Over Time (All Sources)




Figure 3: Attacks over time split by type.

small at 15 attacks. Nonetheless, this is a worthwhile comparison because firms should be reporting to two regulatory bodies, HHS and the SEC. We find that in six cases, the breach is reported to both regulators. However, in seven cases, only HHS is notified. In three cases, the attacks also appear in Hackmageddon. Finally, it is worth noting that no attacks are reported to SEC or Hackmageddon without also appearing at HHS.

4.2. Comparing Attack Types

We next examine whether the types of attacks reported differ by data source. Table 1 shows the breakdown. Data breaches are the most common attack type, accounting for 68% of reported attacks overall. Ransomware comes second at 22%, with other system intrusions accounting for 10%. We do observe some differences between Hackmageddon and the SEC, with other intrusions more prevalent in SEC filings. However, these differences are statistically indistinguishable. As expected, most reports to HHS concern data breaches, as the HITECH act requires disclosure when unauthorized breaches of patient health data occur. Nonetheless, one report concerns ransomware, another significant threat to hospitals in recent years.

Figure 3 studies how the attack type varies by year. The biggest trend here is that ransomware peaked in 2020 and 2021, before receding in 2022. Otherwise, data breaches dominate.

Figure 4 offers another way to compare data sources and attack types, this time with Venn diagrams for SEC and Hackmageddon. We can see that data breaches tend to appear most often in Hackmageddon without being reported to the SEC. Ransomware is more likely to be reported to both the SEC and Hackmageddon (59% of cases appear in both sources). By contrast, the catch-all category of other intrusions tend to appear in only one of the data sources but not both. 11 of 13 reports (85%) appear in just one source.

4.3. Timing of Attack Reports to SEC

The SEC’s motivation for guiding firms to disclose cybersecurity risks in the annual filings is to inform investors




Figure 4: Venn diagrams comparing coverage split by attack category.

better. Should the risk of a cyber attack pose a material risk to a company, investors have a right to be informed of that possibility before an incident occurs.

But does that disclosure of cyber risk actually occur before an attack takes place? To answer this question, we examined all 10-K filings that mention the word “cyber” for firms that experienced an attack. Of the 60 attacked firms in our SEC sample, 51 discussed cyber risks *before* they disclosed an attack in a subsequent filing. Hence, there is substantial but incomplete disclosure of cyber risks among attacked firms.

What happens to the disclosures after attacks? In all 60 cases, firms continue discussing the risk of cyber attack in the regulatory filings after an attack occurs. This is not surprising, given that the risk of cyber attacks was actualized in all firms.

5. Conclusion

To truly understand cyber attack behavior, a reliable empirical record is required. This record is emerging through a mixture of volunteer-led efforts and regulatory mandates to disclose attacks in certain circumstances. In this paper, we have investigated three sources of cyber attacks on enterprises – two mandatory, one voluntary – and found that significant gaps in coverage remain.

We confirmed that relying on news reports, as measured through the volunteer-led Hackmageddon initiative, leads to significant underreporting. We also found gaps in the regulatory filings. Interestingly, not all breach announcements that hit the news are mentioned in the subsequent filings to the SEC, despite the SEC’s guidance to do so. For now, we do not see statistically significant differences in the types of reports that are observed by different data sources. We do observe that ransomware tends to appear in multiple sources, while other system intrusions that are not ransomware or data breaches usually show up in only one source. We also observed that when comparing regulatory mandates, health authorities report to HHS but often not the SEC.

In conclusion, we observe that the reporting does contain lots of useful information to improve understanding of attack prevalence. However, the unstructured nature of the reporting itself means that constructing usable datasets requires significant work on the part of researchers to standardize non-standard reports.

The nature of the data collection process has inherent limitations. For the SEC data, we rely on companies to self-report incidents, and they get to decide when an incident is

serious enough to warrant disclosure. Inevitably, this means that less impactful incidents may go unreported there, even if they could be picked up in news outlets such as those tracked by Hackmageddon. We may also miss some incidents since we only examine filings that mention the word “cyber”, though we expect that number to be quite small since the term is in such widespread use.

Moving forward, we plan to develop an automated process to analyze SEC filings to identify when cyber attacks are reported. We also would like to examine other data sources, both public and private, for completeness. Finally, we would like to investigate the timing of regulatory disclosures more closely. For attacked firms, do they accurately predict the kinds of attacks they actually experience later? Is compliance with SEC disclosure guidance increasing over time and how does it vary by sector? These are some of the intriguing questions we would like to investigate further.

Acknowledgements

We gratefully acknowledge support from the US National Science Foundation Award No. 2147505. As part of the open-report model followed by the Workshop on Attackers & CyberCrime Operations (WACCO), all the reviews for this paper are publicly available at <https://github.com/wacco-workshop/WACCO/tree/main/WACCO-2023>.

References

- [1] R. Anderson and T. Moore, “The economics of information security,” *Science*, vol. 314, no. 5799, pp. 610–613, Oct. 2006.
- [2] T. Moore, “The economics of cybersecurity: Principles and policy options,” *International Journal of Critical Infrastructure Protection*, vol. 3, no. 3–4, pp. 103 – 117, 2010. [Online]. Available: <https://tylermoore.utulsa.edu/ijcip10.pdf>
- [3] Privacy Rights Clearinghouse, “Data breach chronology,” 2023, <https://privacyrights.org/data-breaches>.
- [4] U.S. Department of Health and Human Services Office for Civil Rights, “Breach portal: Notice to the secretary of HHS breach of unsecured protected health information,” 2023, https://ocrportal.hhs.gov/ocr/breach/breach_report.jsf.
- [5] U.S. Securities and Exchange Commission, “Cf disclosure guidance: Topic no. 2,” 2011, <https://www.sec.gov/rules/other/2018/34-83723.pdf>.
- [6] —, “Commission statement and guidance on public company cybersecurity disclosures,” 2018, <https://www.sec.gov/rules/interp/2018/33-10459.pdf>.

- [7] —, “Cybersecurity risk management, strategy, governance, and incident disclosure,” 2022, <https://www.sec.gov/rules/proposed/2022/33-11038.pdf>.
- [8] P. Passeri, “Hackmageddon,” 2023, <https://www.hackmageddon.com/>.
- [9] Temple University, “Critical infrastructure ransomware attacks (CIRA),” 2023, <https://sites.temple.edu/care/cira/>.
- [10] K. Campbell, L. A. Gordon, M. P. Loeb, and L. Zhou, “The economic cost of publicly announced information security breaches: empirical evidence from the stock market,” *Journal of Computer Security*, vol. 11, no. 3, pp. 431–448, Jan. 2003, publisher: IOS Press. [Online]. Available: <https://content.iospress.com/articles/journal-of-computer-security/jcs192>
- [11] A. Acquisti, A. Friedman, and R. Telang, “IS THERE A COST TO PRIVACY BREACHES? AN EVENT STUDY,” in *Workshop on the Economics of Information Security*, 2006.
- [12] S. Gay, “Strategic news bundling and privacy breach disclosures,” *Journal of Cybersecurity*, vol. 3, no. 2, pp. 91–108, Jun. 2017. [Online]. Available: <https://doi.org/10.1093/cybsec/tyx009>
- [13] W. W. Li, A. C. M. Leung, and W. T. Yue, “Where is IT in Information Security? The Interrelationship Among IT Investment, Security Awareness, and Data Breaches,” *MIS Quarterly*, vol. 47, no. 1, pp. 317–342, 2023.
- [14] T. Moore and R. Clayton, “The consequence of non-cooperation in the fight against phishing,” in *Third APWG eCrime Researchers Summit*, Atlanta, GA, October 2008.
- [15] A. Pitsillidis, K. Levchenko, C. Kreibich, C. Kanich, G. Voelker, V. Paxson, N. Weaver, and S. Savage, “Botnet Judo: Fighting Spam with Itself,” in *Proceedings of the 17th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, USA, March 2010.
- [16] V. G. Li, M. Dunn, P. Pearce, D. McCoy, G. M. Voelker, and S. Savage, “Reading the Tea leaves: A Comparative Analysis of Threat Intelligence,” 2019, pp. 851–867. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/li>
- [17] X. Bouwman, H. Griffioen, J. Egbers, C. Doerr, B. Klievink, and M. van Eeten, “A different cup of TI? The added value of commercial threat intelligence,” in *USENIX Security Symposium*, 2020.