# Quantifying and Analyzing Information Security Risk from Incident Data

Gaute Wangen

Norwegian University of Science and Technology,
Teknologiveien 22, 2802 Gjøvik, Norway
gaute.wangen@ntnu.no

**Abstract.** Several cyber security risk assessment and root cause analysis methods propose incident data as a source of information. However, it is not a straightforward matter to apply incident data for risk management. This paper builds on previous work on incident classifications and proposes a method for quantifying and risk analyzing incident data for improving decision-making. The approach was developed using a set of incident data to derive the causes, outcomes, and frequencies of risk events. The data in this paper was gathered from a year of incident handling from a Scandinavian university's security operations centre (SOC) and consists of 550 handled incidents from November 2016 to October 2017. By applying the proposed method, this paper offers an empirical insight into the risk frequencies of the University during the period. We demonstrate the utility of the approach by deducting the properties of the most frequent risks and creating graphical representations of risks. Our primary contribution is the method on how to obtain frequencies and probabilities from incident data and the insight gained from them in risk analysis. This study only defines adverse outcomes and does not include consequence estimates.

**Keywords:** Information security · Cyber security · Security Incidents · Risk Analysis · Threat Intelligence

## 1 Introduction

The topic of this paper is how to categorize, quantify, and apply an organization's information security (InfoSec) incident register for risk analysis.An InfoSec incident is an event or occurrence that contains a breach of either confidentiality, integrity or availability of information or a service. An incident is typically handled by the computer security incident response team (CSRIT) or Security Operations Centre (SOC). Such teams usually consists of InfoSec experts who aim to resolve the event and return the system to normal operations. The SOC often uses an incident management systems which contain records of the incident and steps taken to resolve the incident. The incident handling process goes as follows: an incident gets reported to or detected by the SOC and is assigned to an incident handler (IH). The IH determines if it is an incident and if it belongs to the SOC. If so, he attempts to resolve the incident and records actions

taken. This makes incident data a readily available data source in many organizations for improving the InfoSec risk analysis (ISRA). Quantification is the act of counting and measuring observations into quantities. The risk analysis conducted in this study consists of identifying and quantifying issues that contribute to a risk and analyzing their significance. Previous work has gone into the analysis of InfoSec incidents and their cost [7, 6, 13, 15], but the literature is scarce on the process of how to adapt specific incident data into ISRA. The problem can be described by the follow sequence of events: An phishing e-mail arrives in an employee inbox. The e-mail contains a malicious attachment. An employee opens the attachment and a malware trojan infects the computer and connects back to the attacker. The attacker uploads a keylogger to the infected computer to extract the company username and password. The attacker later uses the stolen credentials to log into the company network to look for vulnerable servers from which he can exfiltrate information. The industry practice is to classify an incident under one specific category as proposed by FIRST[1], US-CERT [2], and others [13, 1, 3, 6, 5]. However, how does one classify the described incident? Is it a social engineering attack, malware infection, compromised user, or data loss? The inherent ambiguity of incident classifications is not sufficiently addressed in current research and there is more useful data that can be used for knowledge gathering and risk analysis in the incident registers.

There are additional obstacles to quantifying incident data which we address in this study: Firstly, no two incidents are identical. To be able to quantify the incidents, we first need to categorize them in a meaningful manner. There already exists classification frameworks, such as those proposed by FIRST and CERT-US, together with taxonomies of computer security incidents (e.g. [11, 8]) that provide a starting point. The desirable level of granularity must be decided by the risk analyst, but for this study we operate with two levels of granularity for the classifications. A computer incident consists of more information that can be quantified for decision-making than can be described by a one incident classification (e.g. "Data leakage"). The previously mentioned classification schemes do not sufficiently recognize the usefulness of adapting more parts of the incident data into risk analysis. Although a typical security incident consists of more than two steps, our data analysis revealed that for the majority of cases we could deduce both a cause and an outcome. Where the former relates to the threat, attack vector and vulnerability and the latter relates to the malicious action taken, intent and outcome. In terms of security controls, the former relates to preventive barriers and the latter relates to reactive barriers.

Through a study conducted at the SOC in the Norwegian University of Science and Technology (NTNU) we have gathered data from 550 incidents. The main purpose of this paper is to show how the risk analyst can categorize and quantify

[1] Forum of Incident Response and Security Teams `https://www.first.org/resources/guides/csirt_case_classification.html`(Visited May 2019)

[2] Federal Incident Reporting Guidelines `https://www.us-cert.gov/government-users/reporting-requirements`(Visited May 2019)

incident data into cause and outcomes, and apply in risk analysis. Specifically, we address the following research questions: (i) How can incident data be quantified and applied in risk analysis? (ii) What does the risk picture look like at the University based on the incident data? and (iii) How can incident data be applied to graphical risk analysis?

We address research question (i) by proposing a risk classification and quantification scheme. Research question (ii) is addressed by analyzing the quantified data from the case study. The final research question (iii) is answered by applying the quantified incident data in different risk analysis schemes to demonstrate the utility and extract knowledge about a specific set of risks.

The remainder of the paper is organized as follows, Section 2 provides general background information on incident classification and quantitative risk assessment. In section 3, we describe how the framework was developed and applied. Furthermore, we describe the data collection process and the applied classifications and risk analysis. Section 4 presents the study with the risk picture for the institution from applying the proposed method. Furthermore, this paper extends the risk analysis in Section 5 where specific risks are studied as examples. Section 6 evaluates the proposed method including the limitations of the study and the proposals for future work. Lastly, we conclude the work in section 7.

## 2 Background and related work

Firstly, this section presents the preliminary work and relevant reports applied in this study. Furthermore, we address the previous work on existing InfoSec incident classification. Lastly, we survey the relevant literature on risk quantification for InfoSec.

This paper builds on the preliminary work in incident classification for root cause analysis published by Hellesen et al. [9] and use of the critical incident tool. Additionally, Chapman [5] has published a policy notice on cyber-security in higher education where he outlines key security challenges faced by the UK higher education and research. Chapman also includes incident statistics from the UK based Janet network in the period January - December 2018. This paper also applies results from the technical report "Unrecorded security incidents at NTNU 2018" [17] which contains statistics from a security awareness survey conducted at the University. Both the statistics from the Janet network and the technical report are used for comparison in this paper.

Kjaerland proposes a taxonomy of computer security incidents based on *Method of operation* and *Impact* which recognizes the attack and the impact as components of the incident [11]. The Method of operation category consists of malicious actions and attack vectors an attacker can apply. The Impact variable consists of the attacker's intentions, such as disrupt and destruct. Kjaerland's approach provides a starting point for incident classifications. Hansman and Hunt [8] proposes a technical taxonomy for incident classification, containing four dimensions or

six levels of categorization per incident. The information gathered for this taxonomy is useful for deep analysis of each incident, but the proposed level of technical detail is typically not needed for risk quantification. ENISA's *Reference Incident Classification Taxonomy* [3] provides the basic categories for the proposed framework. While the *The Common Taxonomy for Law Enforcement and The National Network of CSIRTs* [1] was published by ENISA and Interpol to bridge the gap between the CSIRTs and international Law Enforcement communities. It adds a legislative framework to facilitate the harmonization of incident reporting to competent authorities, the development of useful statistics and sharing information within the entire cybercrime ecosystem. It proposes nine broad categories for incident classification with sub-classifications based on malicious actions. Common for all of these approaches is that they are not developed specifically for risk classification. However, they provide a solid starting point for an incident classification framework scoped for risk assessment.

One of the primary goals of this paper is to quantify frequencies of occurrences for InfoSec risks. Recently there has been multiple attempts at quantifying the cost of information risk incidents. The trend in loss quantification has been for security vendors and other parties responsible for surveys to publish loss estimates. Florencio and Herley [7] discuss the weaknesses such cyber-crime surveys and how the results can entail large amounts of uncertainty. Edwards et.al. [6] investigate a similar problem confined to reported data record breaches from 2006 to 2015. The authors focus on the likelihood component and demonstrate how to derive estimates and predictions about data breaches. The analyzed breaches in the study are divided into negligent and malicious breaches with eight subcategories. Similar to the study in this paper, Kuypers et.al. [13] tackles the problem of incident classification and analysis using eight categories for incident classification. Kuypers also divides the incidents into severity based on the time spent on handling each incidents. The study utilizes 60,000 incident records collected over a 6 year period. The scope of the Kuypers et.al. paper is primarily quantification for predicting trends of frequencies of events and losses. This scope differs from the current paper in that the we conduct deeper analysis by investigating the threat action and applying typical InfoSec risk analysis methods [18] to the data.

Bernsmed et al. [4] have published an illustrative paper for applying bow-tie analysis in cybersecurity. Although, a slightly different approach than Bernsmed, we found the best-suited risk approach for the incident data was the *Bow-tie* as it utilizes both causes and outcomes of each risk or incident. Furthermore, this study builds on Wangen et al.'s [18] Core Unified Risk Framework (CURF), which is a bottom-up classification of ISRA methods and contains an extensive overview of ISRA method content.

## 3 Method

This section outlines the applied method for the study, starting with classification framework development. Furthermore, this section describes data collection, risk analysis and statistics.

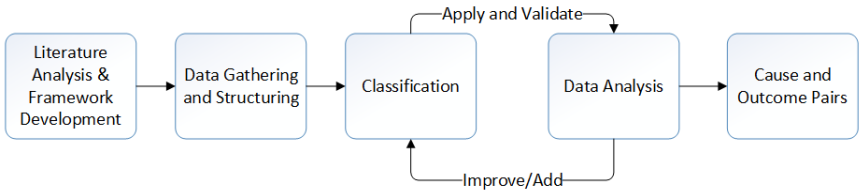


Fig. 1: Development scheme for Incident classification framework

### 3.1 The Incident Classification and Analysis Scheme

As mentioned in the related work, there exist multiple frameworks for classifying computer security incidents, and the proposed classification scheme builds on these as a starting point. There are some underlying premises to the framework: An incident must have (at least) one cause and one outcome. Where the cause relates to a threat exploiting a vulnerability. The outcome is the empirically observable malicious action taken by the threat where he acts on objectives. No two incidents are identical, which means that we must classify to quantify. Specifically, the framework idea is as follows: Start with a set of level 1 incident categories tailored to the organization. Furthermore, the risk analyst identifies a cause and an outcome for each reviewed incident and quantifies them. If either the identified cause or outcome is not present in the current set, the analyst either adds it to the classification set or embeds it into an existing category. By applying this approach, the classification framework receives continuous validation and improvement, visualized in Fig. 1. The classification scheme applied in this study consists of fourteen main classifications with sub-classification. For simple incident classification it might be sufficient to apply the level 1 categories (e.g. [13, 6]). For the dataset in this study, the applied incident classifications are listed in Table 2. The incident analysis (Table 2) shows that some incident categories are primarily applied for cause classifications, others are primarily outcome related, and some are overlapping. For practical reasons we have chosen to define both the cause and outcome within the same categories.

For the data collected in this study, the paper trail of an incident consists of the original incident trigger, the IH's log of steps taken to resolve the incident, and all correspondence with affected parties. In some of the incidents, there might be uncertainty regarding the initial cause, for these cases, we solved this problem by adding broader categories, such as "System compromise," to quantify the

incidents where we have limited knowledge. A limitation of the incident data is that the record sometimes does not include the real cause. For example, a system might be compromised with a Trojan and the outcome is data exfiltration. However, the cause for the initial infection may be unknown. This is also why that in some cases, for example, a compromised user can be the cause of incident and in other cases it might be the outcome. The level 1 categories we have applied for classification are described in Table 1.

### 3.2 Data Collection

The NTNU SOC was established during 2016 and was building capabilities during the time of data collection. All the data presented in this study was extracted from the incident management system at NTNU's SOC, "Request Tracker for Incident Response" (RTIR) from Best Practical Solutions. The dataset presented in this study includes all incidents handled by the SOC between Nov 2016 and Oct 2017, 550 incidents in total. which includes incidents triggered by in-house capabilities, user reported, and third-party reports. The latter include for example other computer emergency response teams notifications, users, vulnerability reports. At the time of data collection, the NTNU SOC consisted of 4 dedicated incident handlers and one part-time member dedicated to working with email-related issues. We qualitatively analyzed each of the incident reports and assigned a cause and an outcome within the classification scheme described previously.

### 3.3 Risk Analysis and Statistics

The ISRA approach in this paper primarily builds on the ISO/IEC 27005:2011 (ISO27005) [2] for understanding risk management. A risk in our proposed ISRA consists of a scenario with an adverse outcome, with a probability distribution of consequences. The ISO27005 defines the scenario as a combination of assets, vulnerability, threat, controls, and outcome. The analysis in this paper quantifies causes and outcomes and determines the frequency of occurrence for each separately, and cause and outcome pairs together. The study applied IBM SPSS for descriptive statistics. The causes and outcomes are analyzed using descriptive statistics, histograms, time-series, and cause-effect flow charts. We apply confidence intervals (CI) for examples of event prediction where we apply the mean and 90% CI as proposed for risk analysis by Seiersen and Hubbard [10]. As proposed in the preliminary study by Hellesen et al. [9] the incident data has utility in obtaining knowledge about causes of unwanted outcomes and vice verse. We apply bow-tie diagrams to illustrate the utility of the data set. Bow-tie is a visual representation the "attack flow" illustrating the causes, preventive controls, reactive controls, and outcomes, see Figure 6. The bow-tie diagrams allow for utilization of both cause and outcome for each incident, thus enabling more in-depth analysis of each risk. The bow-tie analysis is a representation of the attack flow starting on the left with an attack. The diagram is then used to illustrate the security controls in place to prevent the unwanted incident, which

Table 1: Incident classification descriptions applied in the study.

| No. | Level 1 | Description |
|---|---|---|
| 1 | Abuse | Abuse refers to the improper or wrongful use of company assets. Which includes spamming using company resources.It can also be hosting illegal content on the company network, misusing access rights granted, or users complaining about abuse. |
| 2 | Unlawful activity | Refers to any activity that is deemed illegal either by law and legislation. This category also includes police petitions on data extradition. |
| 3 | Malware | The malware category is broad and contains multiple sub-classifications of different malware categories. As there are many strains of malware, the sub-classification is limited to address our incident and risk analysis needs. |
| 4 | Reconnaissance | The reconnaissance category relates to incidents triggered by typical adversarial information gathering activities,including network scanning and packet sniffing. |
| 5 | Compromised Asset | A compromised asset refers to a company asset that has been breached and is under adversarial control. The classification scheme applies four sub-categories. An asset or system in this setting refers primarily to servers, computers, smartphones, and tablets. Network device refers to network infrastructures, such as routers, printers, raspberry pies, and other networked devices. While application compromise refers to the breach of a specific application. This category also includes hardware theft. |
| 6 | Compromised User | A compromised user is when the username and password of an account get compromised. The level 2 category separates between admin and regular users based on the difference in consequence, where the former constitutes a more severe breach. |
| 7 | Compromised Information | This category is used for incidents triggered by observed adversarial actions, such as leaking sensitive data, modifying or changing information, unauthorized access, and privacy violations. |
| 8 | Vulnerable Asset | A vulnerable asset is an organizational property that is vulnerable to external and internal attacks. Typically, software or system can be vulnerable due to a novel vulnerability or lack of patch management. Alternatively, there can be a misconfiguration that leaves the asset vulnerable, or it can also be an inherent vulnerability in a protocol or similar that leaves the asset open for abuse. |
| 9 | Denial of Service | Denial of service (DoS) occurs when a service or asset becomes unavailable. A DoS can be distributed (DDOS) from many compromised systems to route traffic to the target or can occur just from a single machine. A DoS can be caused intentionally by an attacker, or there might be an outage or another failure that causes it. One of the tougher issues to tackle in the classification is whether participation in an outgoing DDoS from a vulnerable asset (e.g., reflexive attack). |
| 10 | Social Engineering | Typical social engineering attacks are phishing and spear-phishing. Where the former targets organization-wide and the latter targets specific individuals or groups. Whaling and CEO frauds target CEOs, high ranking company officers, and their co-workers.The category also includes less frequent frauds such as support fraud, phone fraud, and SMS fraud. |
| 11 | Intrusion Attempt | Intrusion attempts are when the adversary attempts to penetrate the system using technical or physical means. Typical examples are brute-force attempts on login screens and executing exploits on seemingly vulnerable systems. Other incident triggers can come from sensor alarms (intrusion detection systems), log analysis, honeypots, and other detection technology. |
| 12 | Policy Violation | For the NTNU case data, we have two overarching policies: Information Security Policy and the IT Policy. Each with his management system consisting of standards, rules, and procedures to follow. Violations of any of them can trigger an incident. |
| 13 | Other | The other category is for security incidents that do not classify in any of the above but still needs to be solved by the SOC. |
| 14 | Outcome: Negligible/ Fixed/ Failed Attack | This category is only for classifying outcomes and is necessary for incidents that either has no observable adverse outcome was solved by the SOC, or qualified as a failed attack. E.g a spam campaign launched by an external attacker that quickly gets blocked and handled by the SOC without any employees receiving or opening links. |

occurs if all the preventive controls fail. Furthermore, the mitigating controls in place to reduce the consequence are listed. Again, if these fail an unwanted outcome will occur which constitute an incident.

To further illustrate the utility of the incident data in traditional ISRA, we apply the ISO27005 [2] approach which builds on the asset, threat, and vulnerability scheme for security risks. A lot of the focus in the InfoSec industry is on the threat [18, 19], so, we explicitly show how to use the incident data for the threat assessment. For this analysis, the premise is that the cause relates to the threat, the attack method, and vulnerability. The incident data is a historical record which allows us to work with risk analysis: Starting with formulating the risks from cause and outcome pairs, and then decomposing the data into assets, threats and vulnerabilities by examining which assets were compromised, exploited vulnerabilities, and examining the malicious actions taken by the attacker once inside the system. The outcome reveals either the motive, targeted asset, threat actor class, and / or intent. The threat actor is broadly classified and the risk scenario is defined as proposed by Potter's *Practical Threat modeling* [14]. Potter emphasises that a threat is as specific as it needs to be, for example, in most cases of threat modeling it does not make sense to divide threat actors into groups of high granularity. If we are mainly working on defence grouping on motivation as proposed by Potter is generally sufficient for deriving security requirements: "ACTOR does ACTION to ASSET for OUTCOME because MOTIVATION." The threat actor categories proposed by Potter are *Nation state (APT)*, *Organized Crime*, *Insiders*, *Hacktivists*, *Script Kiddies*, and *Others*.

## 4 The InfoSec Risk Picture at the University

The case data together with relevant available statistics were collected from the SOC at the Norwegian University of Science and Technology (NTNU). At the time when this study was conducted, the SOC constituency amounted to about 39 700 students and 6 900 full-time equivalent staff. There were approximately 1500 servers and 15000 managed clients in the network. NTNU has eight faculties and academic curriculum in the natural sciences, social sciences, teacher education, humanities, medicine and health sciences, economics, finance, and administration, as well as architecture and the arts. NTNU also provides state-of-the-art research within multiple technology fields which dictates confidentiality requirements.

The following sections outline the results of applying the classification scheme proposed in this paper starting with the overarching distributions (level 1) of causes and outcomes. Further, we provide the distributions applying the level 2 categorizations with more granularity. Lastly, we derive the most frequent of cause and outcome combinations. For each step, we give examples of how to extract useful information for decision-making from the data.

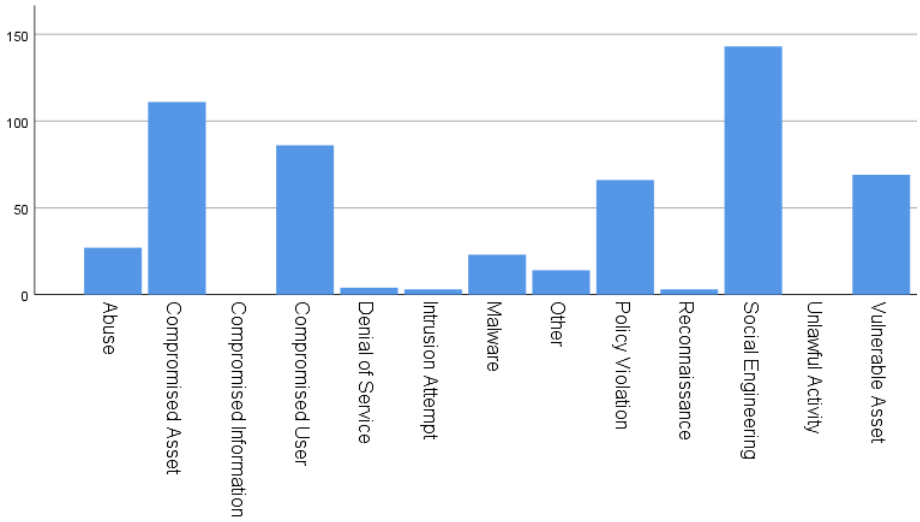


Fig. 2: Incident causes in the NTNU SOC

### 4.1 The Risk Picture According to the Incident Data

The following results start at the top level with the primary causes and outcomes of incidents in their separate histograms. We then describe the level 2 causes and outcomes in a table, before showing the level 1 connection between causes and outcomes. Lastly, we show the trends of causes and outcomes throughout the data collection period.

The total distribution of the causes is illustrated in Fig. 2, and provides an overarching picture of the most common causes of incidents at NTNU. The most common causes in the data set are social engineering attempts (143), compromised assets (107), and compromised users (87). No incidents were caused by actions related to unlawful activity or detection of compromised information.

The total distribution of the 550 outcomes is illustrated in Fig. 2. 201 incidents were handled with a negligible outcome. Furthermore, Abuse (84), Denial of Service (68), and Unlawful activity (55) are the most frequent outcomes of incidents in the constituency.

Table 2 provides the level 2 distribution of both the causes and outcomes for the incidents. A thing to note is that the table only shows quantities and does not contain the connection between the numbers in each column. These numbers provide a higher granularity of information than the level 1 distributions, for example, by decomposing the abuse classification we reveal that the most frequent type of abuse is spam for both causes and outcomes, with misuse of company resources (28) as a common outcome of incidents. The majority of the misuse incidents were caused by a hacking campaign called the *Silent Librarian* by John Chapman[5]. The Silent librarian was launched against Universities and

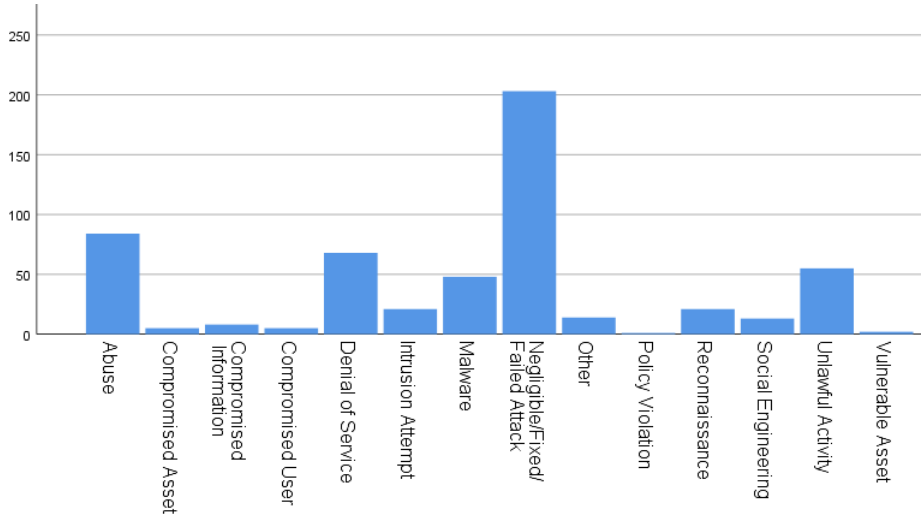


Fig. 3: Incident outcomes in the NTNU SOC

abused the access to publication channels to mass download research articles. In his article, Chapman attributes the campaign to the Iranian state-sponsored group *Mabna* hacking group. The System compromise category is a frequent a cause for incidents (102), while it is seldom an observable outcome of an incident (3). Furthermore, we observe that unlawful activity, reconnaissance, and compromised information are primarily outcome-categories, while the vulnerable asset, compromised asset, social engineering, and policy violations are primarily cause-categories.

Table 3 reveals the level 1 connection between the causes and outcomes in the dataset. The cause is listed on the y-axis and outcome on the X-axis. The table lists the combinations of causes and outcomes for the dataset and describes the frequency of occurrence for each combination. For example, a compromised asset has multiple outcomes, whereas twelve incidents of compromise lead to abuse, thirty-three to malware infections, and sixteen to further reconnaissance. There are several common combinations in the dataset, such as social engineering attempts and fixed attacks (127), which is typically phishing attacks (Table 2). Vulnerable assets often relate to protocol vulnerabilities and similar weaknesses exploited in amplification attacks (29).

### 4.2 Trends and Predictions

Tables 4 and 5 illustrate the development of level 1 causes and outcomes throughout the data collection period. For example, the most frequent cause of incidents in Table 4 is social engineering attempts. These attacks occur regularly each month throughout the dataset with an uptick the three last months, N = 143 with min = 4 and max = 23. The quality of the phishing attacks are often

Table 2: Frequencies of 550 incidents categorised on Cause and Outcome. No link between the Cause and Outcome columns.

| **Classification (Level 1)** | **Sub-Classification (Level 2)** | **Cause** | **Outcome** |
|---|---|---|---|
| Abuse (1) | Spam | 25 | 53 |
| | Illegal Content | | 1 |
| | User Complaint | 1 | 2 |
| | Misuse | | 28 |
| | Web Site copying | 1 | |
| Unlawful activity (2) | Copyright / Piracy | | 55 |
| Malware (3) | Virus | 1 | 3 |
| | Worm | | 1 |
| | Backdoor / Rootkit | | 2 |
| | Trojan | 15 | 31 |
| | Spyware / Adware | 1 | |
| | Hacking tools, Exploits, & Exploit kits | 1 | 1 |
| | Ransomware | 3 | 6 |
| | DNS Hijack | | 2 |
| | Unspecified | 2 | 2 |
| Reconnaissance (4) | Scanning | 3 | 21 |
| Compromised Asset (5) | System Compromise | 102 | 3 |
| | Network Device Compromise | 4 | |
| | Application Compromise | | 1 |
| | Hardware theft | 1 | 1 |
| Compromised User (6) | Admin User compromise | 2 | |
| | Regular User compromise | 84 | 5 |
| Compromised Information (7) | Data leakage | | 3 |
| | Unauthorised Modification | | 2 |
| | Unauthorised Access | | 1 |
| | Privacy Violation | | 2 |
| Vulnerable Asset (8) | Misconfiguration | 14 | |
| | Vulnerable Software | 13 | |
| | Vulnerable System | 2 | |
| | 0-Day Vulnerability | 1 | |
| | Open for abuse | 41 | 2 |
| Denial of Service (9) | DoS/DDoS | 2 | 7 |
| | DoS/DDoS Outgoing | 2 | 61 |
| Social Engineering (10) | Phishing | 112 | 10 |
| | Spear Phishing | 12 | |
| | Whaling / CEO Fraud | 19 | 2 |
| Intrusion Attempt (11) | Brute Force | 2 | 19 |
| | Exploit on non-vulnerable system | 1 | 2 |
| Policy Violation (12) | Information Security Policy | 6 | |
| | IT Policy | 60 | 1 |
| Other (13) | Unclassified | 16 | 16 |
| | Hoax | 1 | |
| | Malware Hosting | | 1 |
| Negligible/Fixed/Failed Attack (14) | | | 203 |
| Sum | | 550 | 550 |

Table 3: Cause (Y-axis) and Outcome (X-axis) combinations between Level 1 categories

| | Outcome | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Cause | Abuse | Unlaw. Act. | Malware | Recon | Comp. Asset | Comp. User | Comp. Info. | Vuln. Asset | DoS | Social Engi. | Intr. Att. | Policy Viola. | Other | Fixed/ Failed |
| Abuse | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 25 |
| Unlawful Act. | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Malware | 0 | 0 | 2 | 0 | 0 | 0 | 2 | 0 | 10 | 0 | 3 | 0 | 0 | 6 |
| Recon. | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 2 |
| Comp. Asset | 12 | 0 | 34 | 18 | 2 | 0 | 1 | 0 | 21 | 3 | 16 | 0 | 1 | 3 |
| Comp. User | 68 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 9 | 2 | 0 | 1 | 3 |
| Comp. Info. | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Vuln. Asset | 0 | 0 | 1 | 1 | 2 | 0 | 2 | 2 | 29 | 0 | 0 | 1 | 1 | 30 |
| DoS | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 |
| Soc. Eng. | 2 | 0 | 9 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 127 |
| Intr. Attempt | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 |
| Policy Vio. | 2 | 54 | 1 | 1 | 1 | 0 | 2 | 0 | 0 | 1 | 0 | 0 | 0 | 4 |
| Other | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 9 | 2 |

attempts of tricking employees to give away their passwords, visit malicious webpages, or open malicious attachments. There is also evidence to suggest that the incidents in the dataset is only the tip of the iceberg and that this is an everyday event at the University [17]. The majority of these attacks originate externally and seem to motivated by financial gain.

The compromised asset class is the second most frequent cause of incidents, N = 111 with min = 3 and max = 19 values. When an asset gets compromised the data shows that the most frequent course of action taken by the attacker is to install Trojan malware (29 occurrences). This gets detected and handled by the SOC when the malware attempts to call home, so we do not know more about the intent in these cases. Other frequent outcomes is the the compromised asset gets exploited in outgoing DDoS attacks (21 occurrences), or used as a stepping stone in either scanning (18) or brute force attacks (14).

Furthermore, the data shows that compromised users is consistent cause of incidents throughout the year averaging seven incidents per month. A frequent course of action taken by the attacker is to abuse the compromised user account to distribute spam e-mail on internal network (42 cases) implying a financial motivation. The data also reveals more malicious attempts of abusing the account for social engineering such as phishing (6 attempts) and whaling/CEO fraud (2 attempts).

Considering the results in the incident outcome table, the data show an increase in abuse cases in the spring semester (Jan - Jun mean = 10 per month) compared to the autumn semester (Nov-Dec and Jul-Oct, mean = 4). Among other things, these abuse cases were very likely caused by the *Silent librarian* campaign, which consisted of a wave of attacks involving compromised user accounts and exploitation of the access given by the University to harvest research articles (26 accounts).

Table 4: Yearly development in incident causes

| | NOV | DEC | JAN | FEB | MAR | APR | MAY | JUN | JUL | AUG | SEP | OCT | Sum |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Abuse | 0 | 2 | 0 | 0 | 3 | 1 | 1 | 1 | 0 | 7 | 3 | 9 | 27 |
| Unlawful Activity | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Malware | 3 | 0 | 3 | 5 | 1 | 2 | 1 | 0 | 0 | 2 | 5 | 1 | 23 |
| Reconnaissance | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 2 | 0 | 3 |
| Compromised Asset | 18 | 7 | 8 | 3 | 12 | 8 | 9 | 6 | 6 | 6 | 9 | 19 | 111 |
| Compromised User | 6 | 5 | 9 | 13 | 12 | 8 | 9 | 9 | 5 | 7 | 1 | 3 | 87 |
| Compromised Information | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Vulnerable Asset | 2 | 5 | 10 | 7 | 5 | 3 | 5 | 8 | 3 | 5 | 5 | 11 | 69 |
| Denial of Service | 1 | 0 | 0 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 4 |
| Social Engineering | 10 | 9 | 4 | 10 | 19 | 11 | 8 | 11 | 6 | 15 | 17 | 23 | 143 |
| Intrusion Attempt | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 3 |
| Policy Violation | 1 | 0 | 21 | 27 | 10 | 0 | 0 | 0 | 0 | 6 | 1 | 0 | 66 |
| Other | 4 | 0 | 1 | 1 | 0 | 1 | 0 | 2 | 1 | 2 | 2 | 0 | 14 |
| Sum | 45 | 28 | 57 | 66 | 64 | 34 | 34 | 37 | 21 | 51 | 46 | 67 | |

Although the dataset in this paper is limited, we can construct basic predictive models using confidence intervals. These models will improve over time with more data. The top 5 risks per year are in Table 6. For example, the most frequent cause of incidents are low consequence phishing attacks: For the coming year, we expect to see between 85 and 114 attacks with a 90% CI. The risk entailing breach to the IT policy was primarily caused by copyright violations, similarly to one of the largest categories in the UK incident data [5]. The risk was found unacceptable and measures were taken in March 2017. The effect of the risk treatment can be seen in Table 4 where the number of policy violation drops after April 2017 and the Unlawful activity classification in Table 5.

Table 5: Yearly development in incident outcomes

| | NOV | DEC | JAN | FEB | MAR | APR | MAY | JUN | JUL | AUG | SEP | OCT | Sum |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Abuse | 6 | 3 | 9 | 15 | 10 | 7 | 9 | 10 | 5 | 5 | 2 | 3 | 84 |
| Unlawful Activity | 0 | 0 | 19 | 27 | 8 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 55 |
| Malware | 13 | 4 | 3 | 3 | 5 | 5 | 2 | 0 | 0 | 4 | 1 | 8 | 48 |
| Reconnaissance | 2 | 2 | 2 | 1 | 5 | 1 | 0 | 0 | 0 | 3 | 2 | 2 | 20 |
| Compromised Asset | 1 | 0 | 2 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 5 |
| Compromised User | 1 | 0 | 1 | 0 | 2 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 5 |
| Compromised Information | 2 | 1 | 1 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 2 | 0 | 8 |
| Vulnerable Asset | 0 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 2 |
| Denial of Service | 6 | 5 | 12 | 7 | 8 | 1 | 8 | 4 | 4 | 5 | 5 | 3 | 68 |
| Social Engineering | 1 | 2 | 0 | 0 | 2 | 2 | 1 | 1 | 2 | 0 | 1 | 1 | 13 |
| Intrusion Attempt | 1 | 1 | 1 | 1 | 2 | 3 | 0 | 0 | 1 | 3 | 2 | 6 | 21 |
| Policy Violation | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 |
| Other | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 4 |
| Negligible/Fixed/Failed | 8 | 10 | 3 | 11 | 21 | 14 | 13 | 20 | 8 | 26 | 27 | 40 | 201 |
| Sum | 42 | 28 | 56 | 66 | 64 | 33 | 34 | 36 | 20 | 50 | 43 | 63 | |

## 5 Risk Analysis and Visualisation

The previous section presented the numbers and a brief analysis of incident occurrences and trends. This section presents a more detailed analysis of the causes and outcomes and how the dataset can reveal data for decision-making. We also present an example of the bow-tie analysis scenario where we model the case of malware infections. Lastly, we model the incidents using the asset, threat, and vulnerability paradigm for security risks and derive the frequency of occurrence.

Table 6: Confidence Intervals of top five risks occurring per year.

| **Cause** | **Outcome** | **Per year** | **90%CI Lower** | **90%CI Upper** |
|---|---|---|---|---|
| Phishing | Negligible/Fixed/Failed Attack | 99 | 85 | 114 |
| Breach to IT Policy | Copyright/Piracy | 52 | 42 | 64 |
| Regular User Compromise | Spamming | 42 | 33 | 53 |
| Open for Abuse | DDoS Outgoing | 29 | 21 | 39 |
| System Compromise | Trojan | 29 | 21 | 39 |

### 5.1 Cause and Outcome analysis

The intention of this section is to illustrate the cause and outcome analysis as an initial step into creating a graphical risk representation. We will start by using compromised accounts as an example. Account compromise is one of the most frequent causes of incidents at the NTNU SOC, Table 2. A compromised account is when a company username and corresponding password gets compromised by attackers. Account comprise has previously lead to costly incidents at NTNU, such as the aforementioned *Silent librarian* campaign caused around 15 incidents (uncertain attribution) and incidents where the network is used as a staging point. therefore, a priority risk to resolve. During the year of data collection, there were 84 incidents recorded caused by regular user compromise, averaging seven incidents per month. The trend is illustrated in Fig. 4 which shows a peak in February 2017 with 13 incidents caused by compromised accounts and a low in September 2017 with only one. Analysing the incidents, we find a distribution of outcomes illustrated in Figure 5. The most frequent outcome of an account compromise is spamming internal users, however, several other outcomes are more severe, with misuse of resources and whaling/CEO fraud attempts bearing the potentially most severe consequences. From an outcome perspective, we can apply this data in the decision process to choose consequence reducing measures to control risk. Looking at the causes of user compromise, Table 3 shows that we only have five incidents where user compromise was the known outcome, where all had been caused by social engineering attacks. Using this approach, we can

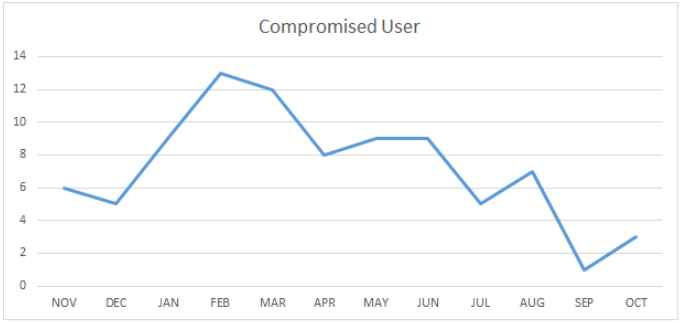


Fig. 4: The figure illustrates the amount of incidents caused by compromised accounts per month.

also reveal areas with uncertainty: for example, the data reveal a lot about the intentions of attackers who use compromised accounts, but the data is lacking on how the accounts are being compromised.

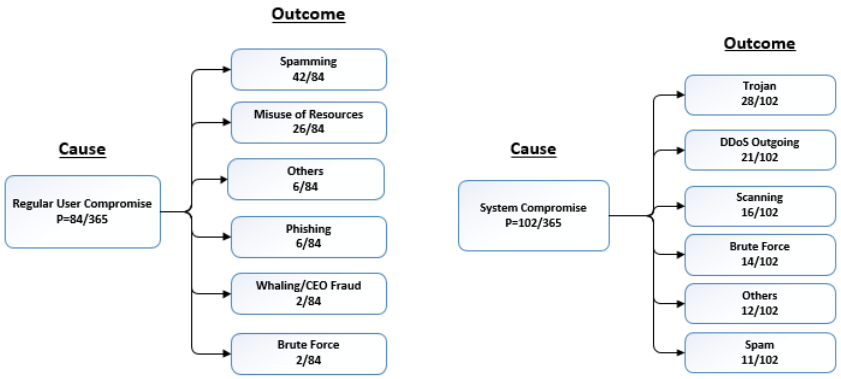


(a) User Compromise outcomes (b) System Compromise outcomes

Fig. 5: Distribution of outcomes from incidents.

Another case with missing initial causes is the system compromise which is the "catch-all" category with 102 occurrences. When attackers compromise systems and establish a foothold in the network, the incident outcomes reveal some of their intent, illustrated in Figure 5. For example, in twenty-eight instances machines are infected with trojans conducting malicious actions. Typically, the observable action is the call home to the Trojan owner. In thirty instances the infected systems are used as staging points for further attacks to compromise more systems through scanning and brute-force. The data also reveals that twenty-

one compromised systems are recruited into botnets and participate in outgoing DDoS attacks on third parties. A thing to note with the system compromise category is that the uncertainty regarding the cause will be reduced with increased detection and forensics capability.

### 5.2 Bow-tie analysis of Malware Infections

This section illustrates the utility of the data for Bow-tie risk analysis (described in Section 3.3). Malware infections are at the root of many severe cybersecurity breaches [16] and are present in the current dataset with twenty-five known causes and forty-six known outcomes. This section contains one attack flow model to illustrate the concept.

To populate the bow-tie model we use the known causes (left side) and outcomes (right side). The unwanted outcome, "Malware infection", is placed in the middle. Further, we have to map out the relevant security controls before we can apply the bow-tie analysis. For the analysis, we are interested in existing preventive controls that reduce the probability of the attack occurring and mitigating controls that reduce the consequence of an incident. Figure 6 illustrates how the classified incident data can be used in bow-tie analysis. All the known causes of malware infections are listed to the left with their known distributions and the outcome distribution to the right. The controls are described at an abstract level to not reveal any defensive capabilities of the SOC. Typically, the bow-tie would make a connection between all causes/outcomes and relevant controls in the figure, but due to the amount and complexity of each incident and the controls involved we created an example instead of including it in the bow-tie.

Figure 7 illustrates how the classified incident data can be used in a bow-tie analysis. The scenario being analyzed is that an attacker succeeds with a phishing attack, infects the victim with malware, and abuses the machine for DDoS. For simplicity, we focus on only one of the attacks paths for the known causes of malware infections listed to the left together with the preventive controls. The controls are listed in the order the attack meets them, e.g. the spam filters will reduce the amount of malicious email that reaches the target a certain amount (the efficiency can be quantified). Furthermore, if the email with the malicious link or attachment reaches the target, the security awareness of the user is his/hers ability to recognize phishing. If this control also fails and the target opens the malicious link or attachment, the end point security is remaining barrier that can prevent an infection. If all controls fail, the malware infection occurs. The relevant consequence mitigating controls are listed to the right together with the undesirable outcome. Typically, the bow-tie would make a connection between all causes/outcomes and relevant controls in the figure, but we have used only one attack patch as an example for the sake of simplicity. The model illustrates all the controls involved both before and after a malware infection, which enables analysis of both preventive and consequence mitigating controls to identify the weaknesses in the security chain. The frequencies of each event aid the decision-maker in determining whether a risk is unacceptable or not and if new controls

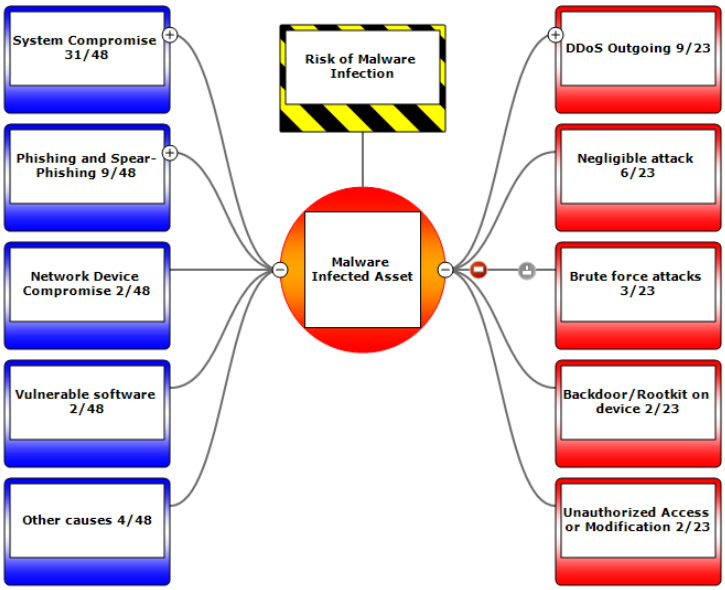


Fig. 6: Bow-tie risk analysis illustrating known causes and outcomes with their respective frequencies for Malware infections in the dataset. Modelled with *BowTieXP*

should be added to the control chain or existing controls should be strengthened. Furthermore, this type of risk analysis will allow for measurements of control efficiency when implementing new barriers in the security chain.

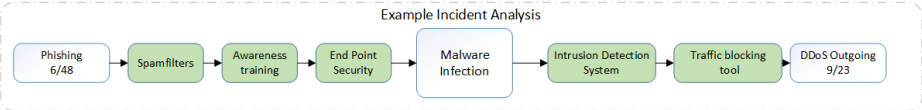


Fig. 7: Example risk analysis of an Incident derived from the bow-tie.

### 5.3 Identifying assets, threats and vulnerabilities

The classic ISRA approach from ISO/IEC 27005:2011 [2] advocates to begin the risk assessment process with asset identification and evaluation, before identifying the threats, controls, and vulnerabilities, and inducing the risk. We have generalized the incident data into the more traditional ISRA [18] in Table 7. The model builds on the categorized versions of the incidents and therefore contains some generalizations. The table provides an organizational risk picture for

decision-making based on incidents. A more detailed model can be obtained through a more thorough analysis of the incident data. From the dataset, the most frequent risk is phishing attacks, while copyright violations are the second most frequent, and user compromise with internal spamming is the third.

Table 7: The five most frequent cause and outcome pairs analysed with an asset, threat, and vulnerability model.

| Nr | Cause | Outcome | Frq. | Target Assets | Threat actors | Motive | Intent | Vulnerability |
|---|---|---|---|---|---|---|---|---|
| 1 | Social Engineering: Phishing | Negligible/Fixed/Failed Attack | 99 | Money, Account credentials, Resources | Criminals | Financial, (intelligence) | Unauthorized access, Misuse, Deny access | Human factor |
| 2 | Breach to IT Policy by registered user | Copyright or Piracy | 52 | Resources (bandwidth) | Insiders | Financial, neglect | Misuse | IT policy |
| 3 | Regular user compromise | Used to send spam | 42 | Money, Account credentials, Resources | Criminals | Financial | Unauthorized access, Misuse | Human factor, weak passwords, etc. |
| 4 | System resource open for abuse | Exploited in outgoing DDoS attacks | 29 | Resources (bandwidth) | Opportunists, Criminals | Revenge, Financial | Deny Access | Weak configuration |
| 5 | Compromised system | Installed Trojans | 28 | Company systems, Secrets | Criminals, APT, Chaotic actors | Financial, Political | Unauthorized access, Deny access, Staging point | Vulnerable systems |

The analysis of the incident data provides strong indication on what the attackers think that the University's primary assets are (see outcomes, Table 2): (i) Computing power and resources - for conducting attacks and as a staging point. This is evident from the amount of attacks launched from network through scanning and brute force attacks . Company accesses are abused to mine resources that are only available through university contracts, such as the Silent librarian campaign, and for the hosting of illegal content. (ii) Bandwidth capacity - recruited in outgoing DDoS attacks and for illegal file sharing in violation of copyright laws. (iii) User and admin accounts - Harvested and traded, gives access to company resources, and is used in phishing/spamming. In addition, financial motives are apparent through attempts of CEO frauds/whaling, phishing, and ransomware. (iv) Information - only seven of the 550 incidents ended in a known information compromise, which indicates that the majority of the attacks that causes incidents aims to exploit other assets at the University. While securing information is important, the data shows that it is the accesses and resources the University governs that were most interesting for the attackers in the time frame. Using risk nr 1 in Table 7, we see that NTNU is also a frequent target of social engineering campaigns. We deduce that the generic motive behind phishing campaigns is financial as they typically target usernames and passwords, financial data, and other resources. This is typical cyber-criminal behaviour using low cost social engineering attacks. The table also reveals that the University network is a popular staging point for launching attacks and recruiting resources into DDoS attacks. We attribute this to opportunists and criminal groups.

Regarding vulnerability, the data reveals two specific weaknesses: social engineering attacks and vulnerable systems. So, to reduce incidents conventional treatments are awareness training and improving the system portfolio and patching routines. Although the dataset is lacking in knowledge about causes for compro-

mised accounts, the primary attack vector against NTNU is phishing attempts. These attacks typically target account information, and although the incident report states an attack was handled, it is likely unrecorded instances of employees falling for the scam and not reporting. In this case, the dataset allows for hypothesis formulations that can be researched in future projects.

The current threat hype in InfoSec is Nation-state backed groups, or so called advanced persistent threat (APT). These groups are typically involved in influence operations, sabotage, and cyber espionage [16]. Due to some of the technological research being done at universities, they are natural targets of espionage looking to gain a technological advantage. Although the incident data do reveal Trojan activity and more advanced attacks, the dataset does not necessarily reveal advanced persistent threat (APT) activity. The skilled actors are better at hiding their tracks and more data is needed to conclude that such actors are present in the systems. This is a limitation of the current dataset which can be addressed with forensic capability.

## 6 Discussion, Limitations, and Future Work

This section discusses the contribution, limitations, and path for future work for each research question proposed in this paper.

### 6.1 Classifying incidents

The contribution of this paper is primarily practical: The proposed classification framework contributes to solving the practical problem of quantifying incident data for risk analysis. As we demonstrated in the study, the proposed method enables an overview of incidents that will improve the understanding of the risk landscape at the organization. Although incident reports may vary in format and content, the proposed framework and method has been validated on 550 incidents and should be adaptable for risk quantification at most organizations having InfoSec incident records. The practical implication of the framework is that it enables simple statistical models of risk frequencies and trends, together with graphical modelling in bow-tie diagrams. The approach also facilitates more complex risk analysis to reduce the uncertainty especially related to frequencies of occurrence and lends itself to the critical incident tool in Root Cause Analysis [9].

All models are simplifications of reality and as mentioned in the introduction, an incident can consist of a chain of causes with multiple adverse outcomes, rather than just one cause and outcome as implied in this framework. We recognize this issue and a more sophisticated approach should be considered in cases where more detailed information is needed. An approach model containing multiple causes, such as a root cause analysis-approach [9] or the Lockheed Martin *Cyber Kill Chain* can be adapted to improve the model. Another limitation is with the method in this paper: while it is likely that the level 1 classifications can be generalized to most organizations and industries, the level 2 classifications

should be tailored to the organization planning to use them. However, both the level 1 and 2 classifications as proposed should provide a starting point for incident classification. A suggestion for future work would be to work on a common framework for a higher detail incident analysis based on traditional ISRA.

Furthermore, all of the categorization done for this paper was done by an analyst, which adds subjectivity in the analysis. We attempted mitigating this issue by applying firm categorizations, outlining rules for categorizations, and adapting the framework as needed. It is a challenge to keep the categories unambiguous and prevent overlap between them. For example, outgoing DDoS is both a DDoS attack and abuse of network infrastructure. Our main categories were developed from the best practice and is similar to those applied in the Jisc SOC [5], this same ambiguity is seen across frameworks. A path for future work is to propose a framework for incident classification using generic risk classifications as a starting point. It is clear from the Table 2 some of the categories are more likely to be a cause of an incident than the outcome and vice verse. Refining causes and outcomes with risk quantification as the goal could assist in improving risk management of cyber security risks.

The incident analysis and classification is a time consuming and repetitive job, which makes automation another path for future work. Depending on the incident record system, the process of automatically classifying incidents could be developed by retrieved the data, produce the dataset for machine learning, develop identifiers for each category, and develop the algorithm for incident classification and risk quantification.

### 6.2 The Risk Picture

The paper also presented the risk picture as seen from incident data at NTNU from Nov 2016 to Oct 2017. The risk picture has limited generalisability, but when comparing the results to the statistics published by Chapman on the security incidents in the UK [5], there is a similarity in the threat landscape across borders in Northern Europe. Compared to the security survey of unreported incidents at the university [17], there is likely under-reporting as only 40% in the survey knew how to report a computer security incident while 48% knew about spear-phishing attempts directed at them or their colleagues. Additionally, 5,2% of the respondents knew about instances where the University infrastructure had been abused for crypto-currency mining, which is not a part of the risk picture according to the incident data. A limitation of this study is therefore that the dataset only contains incidents. We have not included intelligence gathered with other tools, such as intrusion detection systems (IDS), end-point security, surveys, and so forth. A path for future work is to combine data from these sources to produce a more comprehensive risk picture for academia or another industry. Another limitation when it comes to working with historical data is that the records only contain risks that have previously occurred, therefore, basing the risk management purely on historical data will neglect risks that are not in the dataset.

The willingness to share such data in scientific research has been limited [12]. However, publishing data regarding cyber incidents helps to build the theoretical understanding of cyber risks faced by the academic sector and will improve policy. As this paper has highlighted, classifying incident data is not a straight forward matter and it is not always clear what the underlying premises and data source are for generating these statistics: For example, the statistics from the Janet network [5] has a large amount of malware incidents per month. The amount of malware in the incident data is dependent on the anti-malware solution of the SOC and how much trust one puts into each reported "mitigation" from the solution together with how one defines an incident. In this paper, we have tried to be transparent in both where the data came from and how it was treated, such clarity is needed to reduce uncertainty when interpreting the statistics.

### 6.3 Risk Visualization

We applied the bow-tie diagram to illustrate the utility of the dataset. Applying bow-tie diagrams to the data allowed for construction of simple attack flows with frequencies of occurrence both for cause and outcome for each incident category. The strength of this approach is that it is easy to understand and communicate. The risk visualization also allows security control modeling and measurement of control efficiency. The drawback of the diagrams is that they can be an oversimplification of reality in cases of severe risks. In these cases, more sophisticated modelling techniques can be used, such as *event trees* and *attack trees*. Longer attack flow diagrams with multiple causes and outcomes is also a possibility. Another path for improvement is to work on the bow-tie diagrams for the more severe risks and research ways of measuring control efficiency and integrating them into the risk assessment model. This should also include working with loss estimates for the identified outcomes. Adapting Kuypers et.al. [13] approach for differentiating incident consequences based on time spent handling the incident is a start. However, only considering the cost of time spent handling the incident represents a too narrow view on consequences, as there is also possibilities for production loss, asset damage, legal fines, and reputation/competitive advantage to consider. The incident impacts to each of these areas can be estimated by applying the approach proposed by Seiersen and Hubbard [10]. Another limitation with the information presented in this paper regarding asset, threat, and vulnerability are generalizations from the incident data. More specific information on each is present in the incident data but not reported in this paper. A limitation regarding threat actors is that we can not know who they are because of the *attribution problem*[16]. The threat actors are categorized on perceived motivation using the approach proposed by Potter [14] and Wangen et.al. [18, 19].

## 7 Conclusion

This paper has proposed and applied a classification approach for incident data to smooth the transition from incident report to risk quantification and analy-

sis. Our proposed method and framework was anchored in a two-level approach based on established incident classifications and expanded when necessary. The framework was scoped to classify incident causes and outcomes for quantification. By applying the method to a case, we were able to create an empirical risk picture for the University including all the known causes and outcomes of incidents. This study found that one will not get a complete risk picture from analyzing the incident data alone, but it will provide valuable insight into key issues the organization faces: According to the data, the risk picture for NTNU contained a range of cyber attacks, such as social engineering attempts, vulnerable/compromised devices, malware infections, and DDoS. The most frequent threat was identified as social engineering attempts, including phishing, spear phishing, and whaling/CEO frauds. Compromised assets and users made out the second and third most frequent causes of incidents. About 2/5 of the InfoSec incidents were resolved without any observable adverse outcome. Abusing the University infrastructure through outgoing DDoS attacks, spamming, and copyright violations are the three most frequent outcomes in the data set. Although 12 months of data is a short period, we have demonstrated how to apply the proposed approach to study trends in both the cause and the outcome from incidents. Furthermore, the proposed incident analysis method has merit when integrated with the bow-tie analysis, as the model fully utilizes both the cause and outcome statistics. Quantifying incidents allows for predictions, but it also enables the risk analyst to measure the treatment effect over time. In cases such as compromised accounts, we also applied the analysis to detect knowledge gaps whereas the data revealed that there was little knowledge as to how accounts were compromised.

ACKNOWLEDGEMENTS

The NTNU digital security section and SOC consisting of Christoffer Vargtass Hallstensen, Frank Wikstrøm, Harald Hauknes, Hans Åge Marthinsen, Vebjørn Slyngstadli, Gunnar Dørum, Lars Einarsen, and Stian Husemoen. Vivek Agrawal and the anonymous reviewers for help with quality assurance.

## References

1. Common taxonomy for law enforcement and the national network of csirts, version 1.3. Tech. rep., ENISA and Europol E3 (2017), `https://www.europol.europa.eu/publications-documents/common-taxonomy-for-law-enforcement-and-csirts`
2. Information technology, security techniques, information security risk management (ISO/IEC 27005:2011)
3. Reference incident classification taxonomy: Task force status and way forward. Tech. rep., ENISA (January 2018)
4. Bernsmed, K., Frøystad, C., Meland, P.H., Nesheim, D.A., Rødseth, Ø.J.: Visualizing cyber security risks with bow-tie diagrams. In: International Workshop on Graphical Models for Security. pp. 38–56. Springer (2017)
5. Chapman, J.: How safe is your data? cyber-security in higher education. HEPI Policy Note **April**(12) (2019)

6. Edwards, B., Hofmeyr, S., Forrest, S.: Hype and heavy tails: A closer look at data breaches. Journal of Cybersecurity **2**(1), 3–14 (2016)
7. Florêncio, D., Herley, C.: Sex, lies and cyber-crime surveys. In: Economics of information security and privacy III, pp. 35–53. Springer (2013)
8. Hansman, S., Hunt, R.: A taxonomy of network and computer attacks. Computers & Security **24**(1), 31 – 43 (2005)
9. Hellesen, N., Torres, H., Wangen, G.: Empirical case studies of the root-cause analysis method in information security. International Journal On Advances in Security **11**(1&2) (2018)
10. Hubbard, D.W., Seiersen, R.: How to measure anything in cybersecurity risk. John Wiley & Sons (2016)
11. Kjaerland, M.: A taxonomy and comparison of computer security incidents from the commercial and government sectors. Computers & Security **25**(7), 522–538 (2006)
12. Kotulic, A.G., Clark, J.G.: Why there aren't more information security research studies. Information & Management **41**(5), 597–607 (2004). https://doi.org/10.1016/j.im.2003.08.001, `http://dx.doi.org/10.1016/j.im.2003.08.001`
13. Kuypers, M.A., Maillart, T., Pate-Cornell, E.: An empirical analysis of cyber security incidents at a large organization. Department of Management Science and Engineering, Stanford University, School of Information, UC Berkeley **30** (2016)
14. Potter, B.: Practical threat modeling. Login **41**(3) (2016), `https://www.usenix.org/publications/login/fall2016/potter`
15. Romanosky, S.: Examining the costs and causes of cyber incidents. Journal of Cybersecurity **2**(2), 121–135 (2016)
16. Wangen, G.: The role of malware in reported cyber espionage: A review of the impact and mechanism. Information **6**(2), 183–211 (2015)
17. Wangen, G., Brodin, E.Ø., Skari, B.H., Berglind, C.: Unrecorded security incidents at NTNU 2018 (Mørketallsundersøkelsen ved NTNU 2018). NTNU Open Gjøvik (2019)
18. Wangen, G., Hallstensen, C., Snekkenes, E.: A framework for estimating information security risk assessment method completeness. International Journal of Information Security pp. 1–19 (2017)
19. Wangen, G., Shalaginov, A., Hallstensen, C.: Cyber security risk assessment of a ddos attack. In: International Conference on Information Security. pp. 183–202. Springer (2016)